



US 20180181315A1

(19) **United States**

(12) **Patent Application Publication**

KUSTERS et al.

(10) **Pub. No.: US 2018/0181315 A1**

(43) **Pub. Date: Jun. 28, 2018**

(54) **DATA STORAGE SYSTEM WITH MULTIPLE DURABILITY LEVELS**

Publication Classification

(71) Applicant: **Amazon Technologies, Inc.**, Seattle, WA (US)

(51) **Int. Cl.**
G06F 3/06 (2006.01)

(72) Inventors: **NORBERT P. KUSTERS**, REDMOND, WA (US); **NACHIAPPAN ARUMUGAM**, SEATTLE, WA (US); **CHRISTOPHER NATHAN WATSON**, SEATTLE, WA (US); **MARC JOHN BROOKER**, SEATTLE, WA (US); **DAVID R. RICHARDSON**, SEATTLE, WA (US); **DANNY WEI**, SEATTLE, WA (US); **JOHN LUTHER GUTHRIE, II**, SEATTLE, WA (US)

(52) **U.S. Cl.**
CPC **G06F 3/0608** (2013.01); **G06F 3/0619** (2013.01); **G06F 3/0622** (2013.01); **G11B 33/128** (2013.01); **G06F 3/0659** (2013.01); **G06F 3/067** (2013.01); **G06F 3/0644** (2013.01)

(73) Assignee: **Amazon Technologies, Inc.**, Seattle, WA (US)

(57) **ABSTRACT**

(21) Appl. No.: **15/392,857**

A data storage system includes multiple head nodes and multiple data storage sleds mounted in a rack. For a particular volume or volume partition one of the head nodes is designated as a primary head node for the volume or volume partition. The primary head node is configured to store data for the volume in a data storage of the primary head node and cause the data to be replicated to a secondary head node. The primary head node is also configured to cause the data for the volume to be stored in a plurality of respective mass storage devices each in different ones of the plurality of data storage sleds of the data storage system.

(22) Filed: **Dec. 28, 2016**

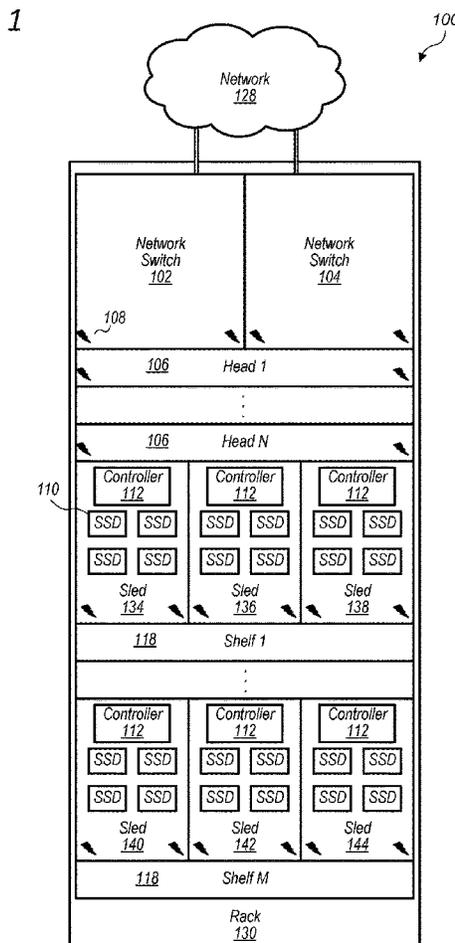
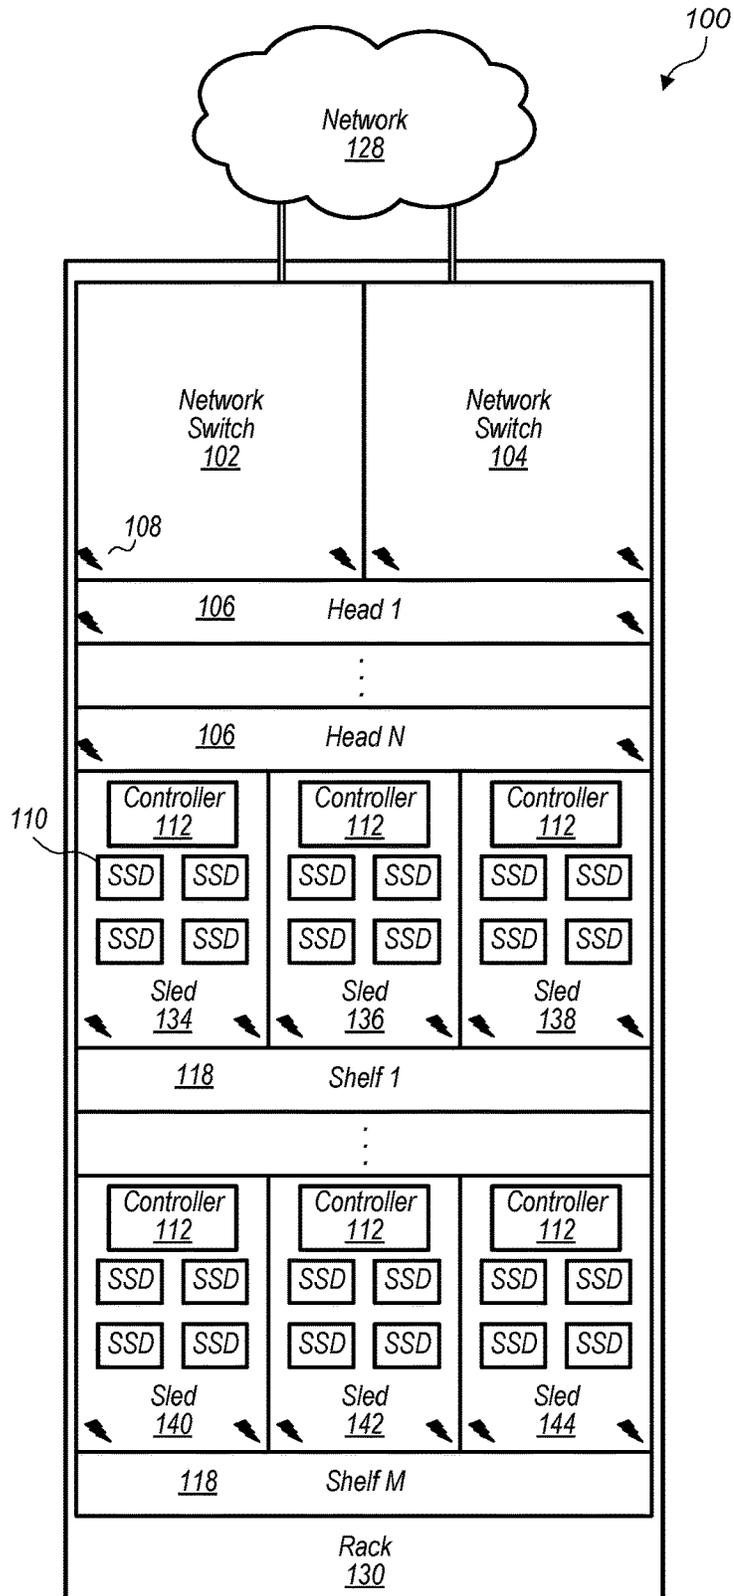


Fig. 1



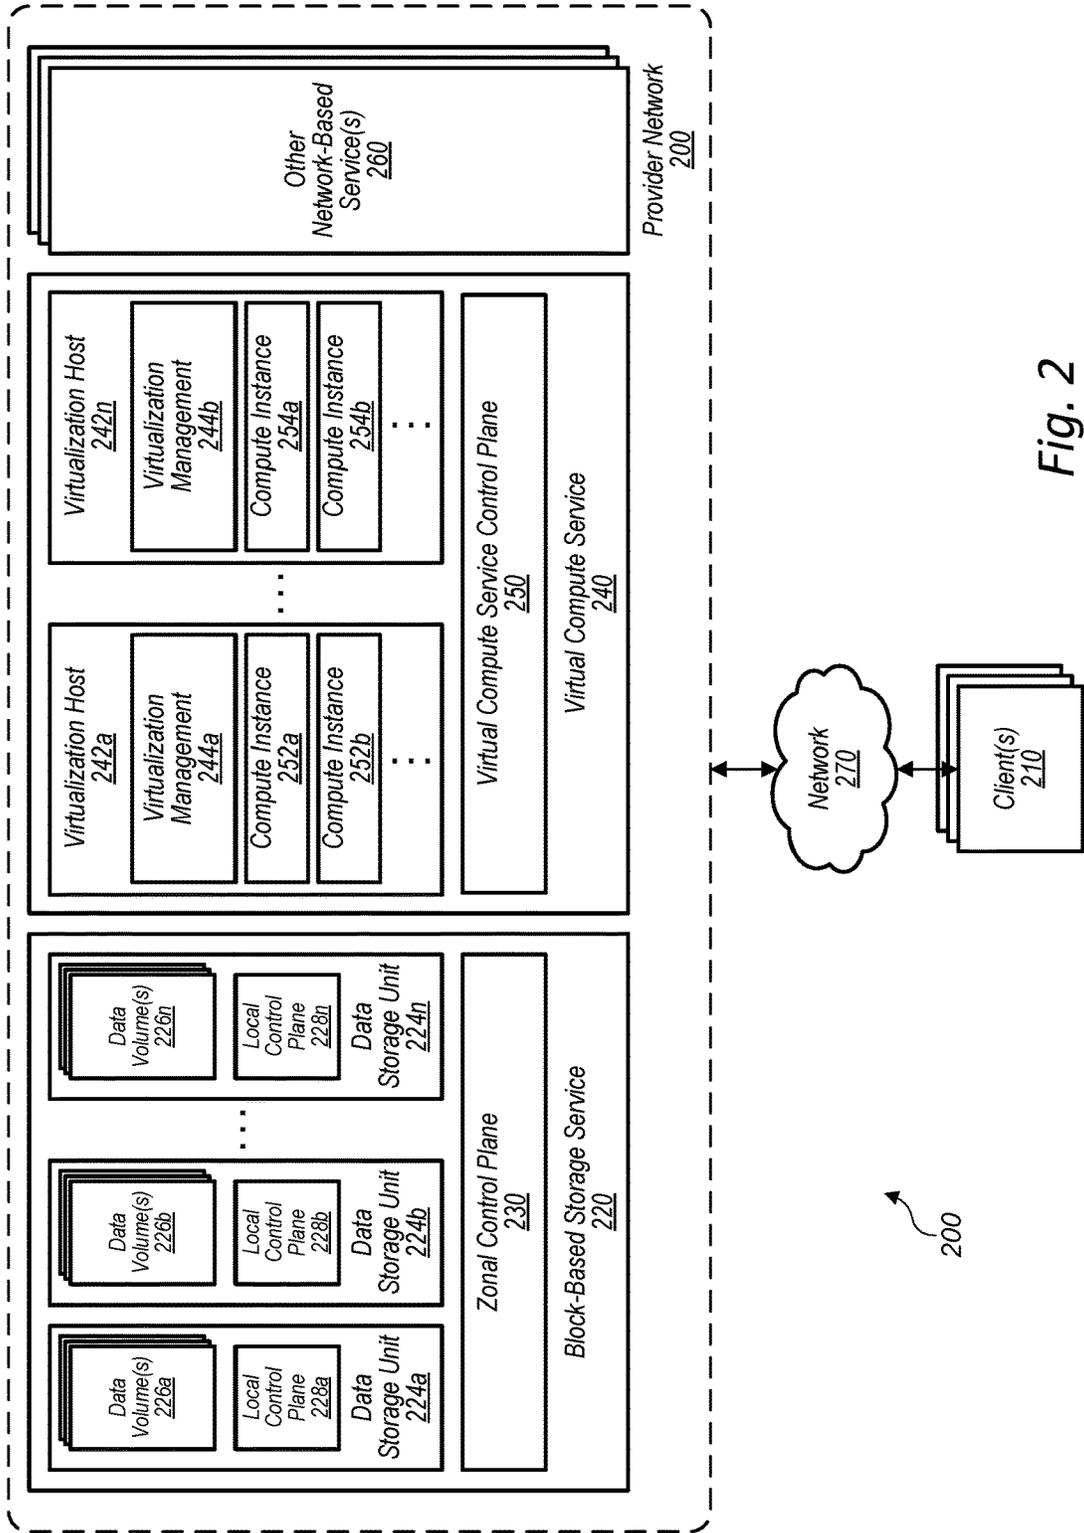


Fig. 2

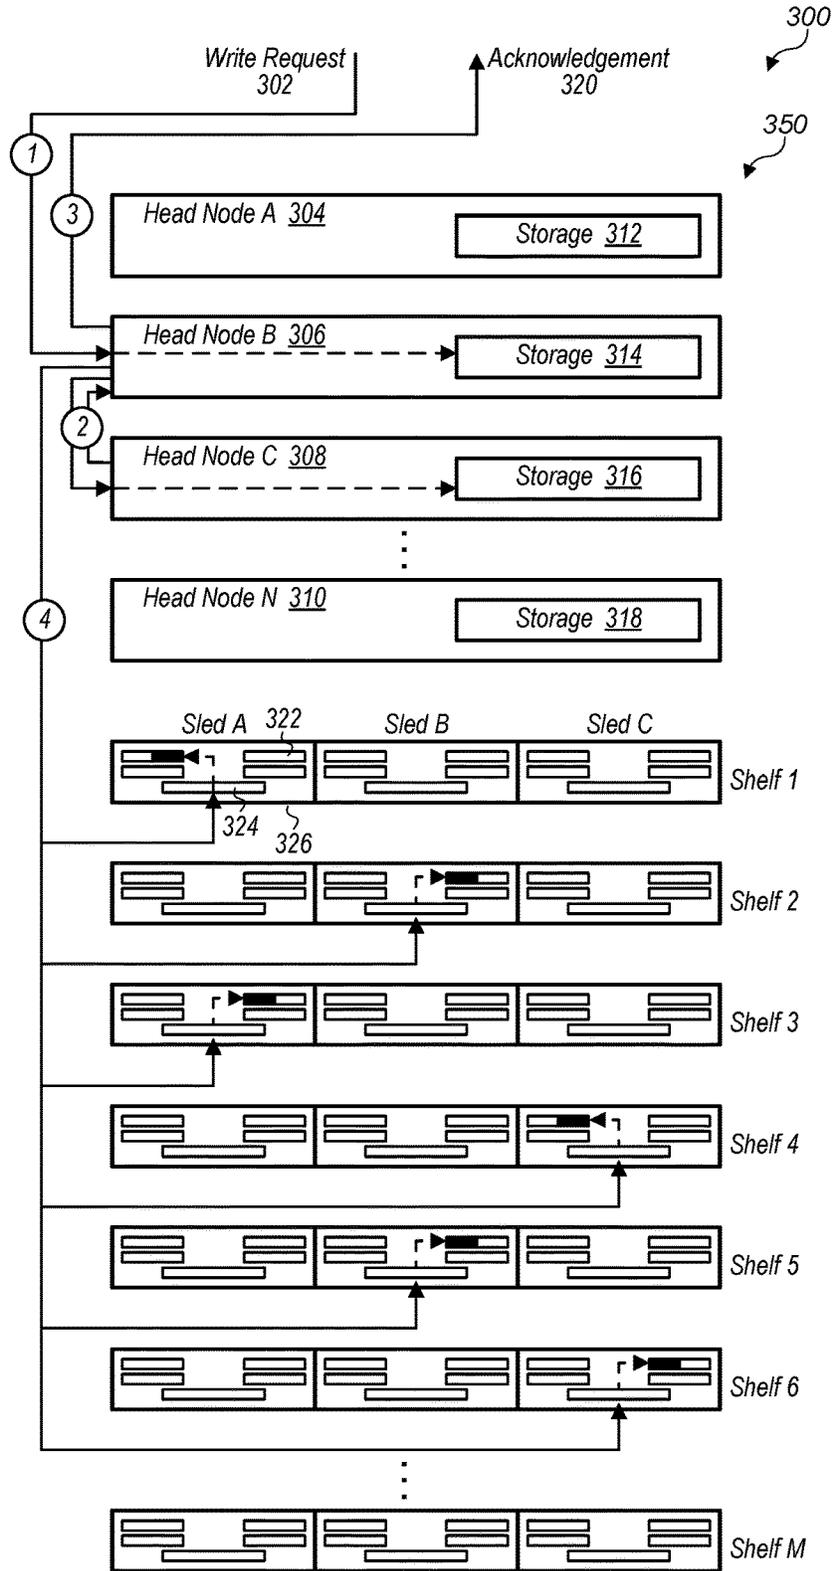


Fig. 3

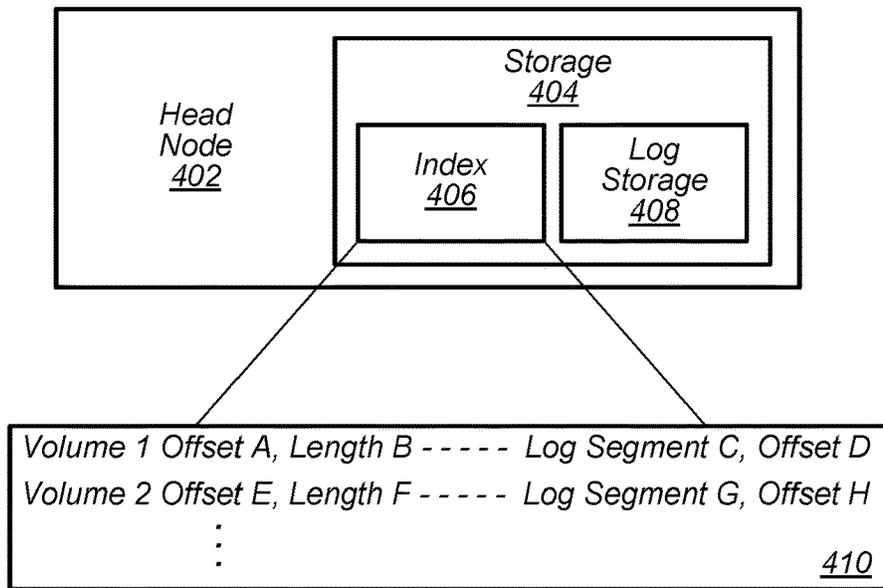


Fig. 4A

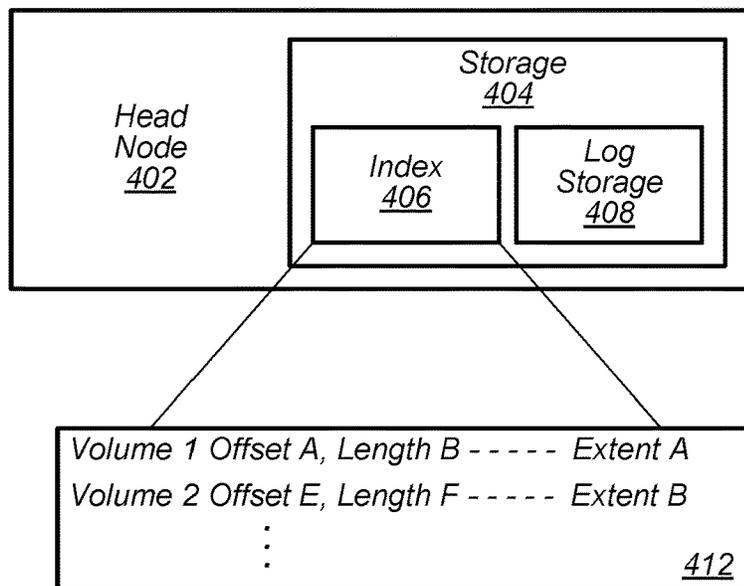


Fig. 4B

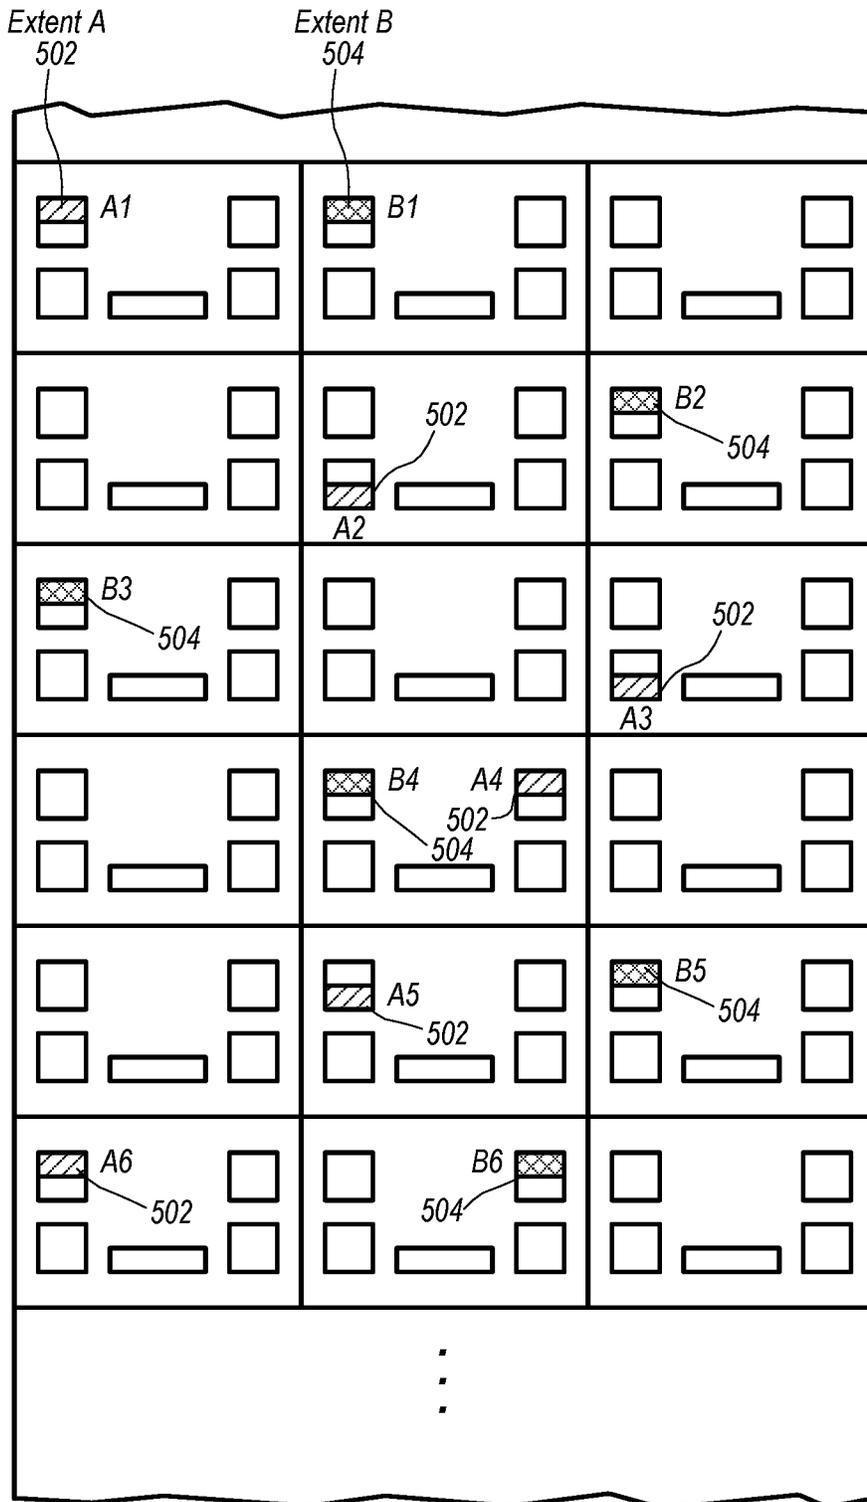


Fig. 5

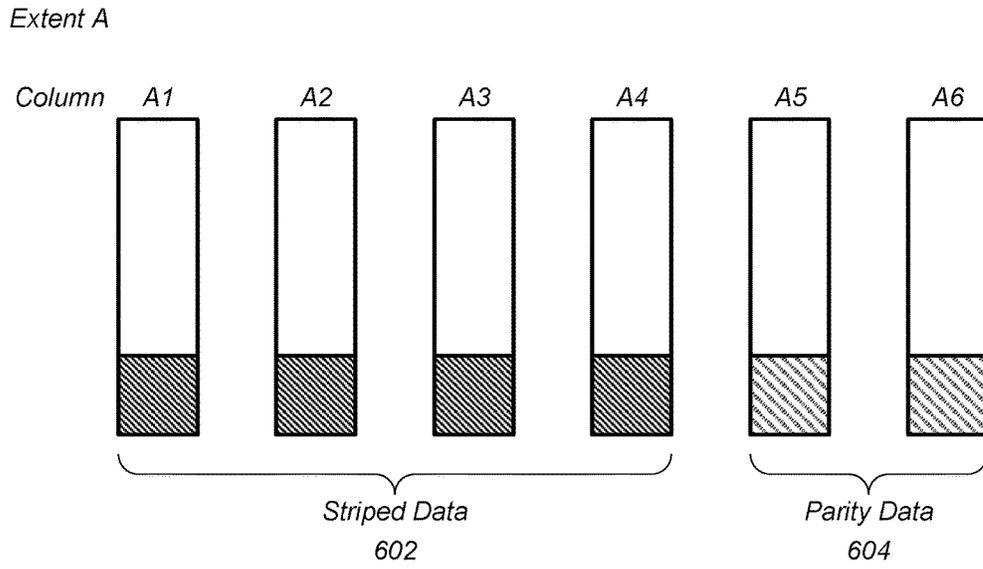


Fig. 6A

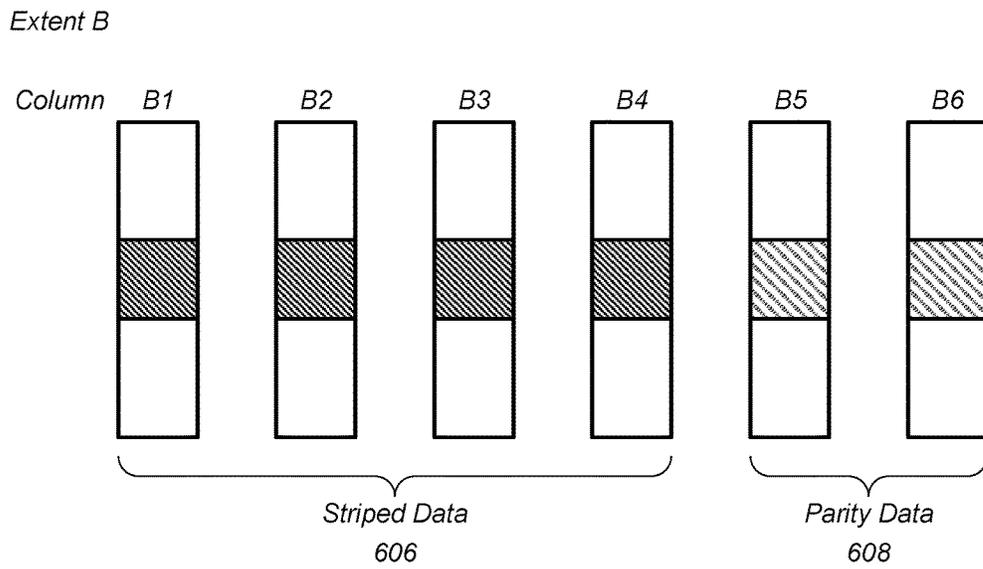


Fig. 6B

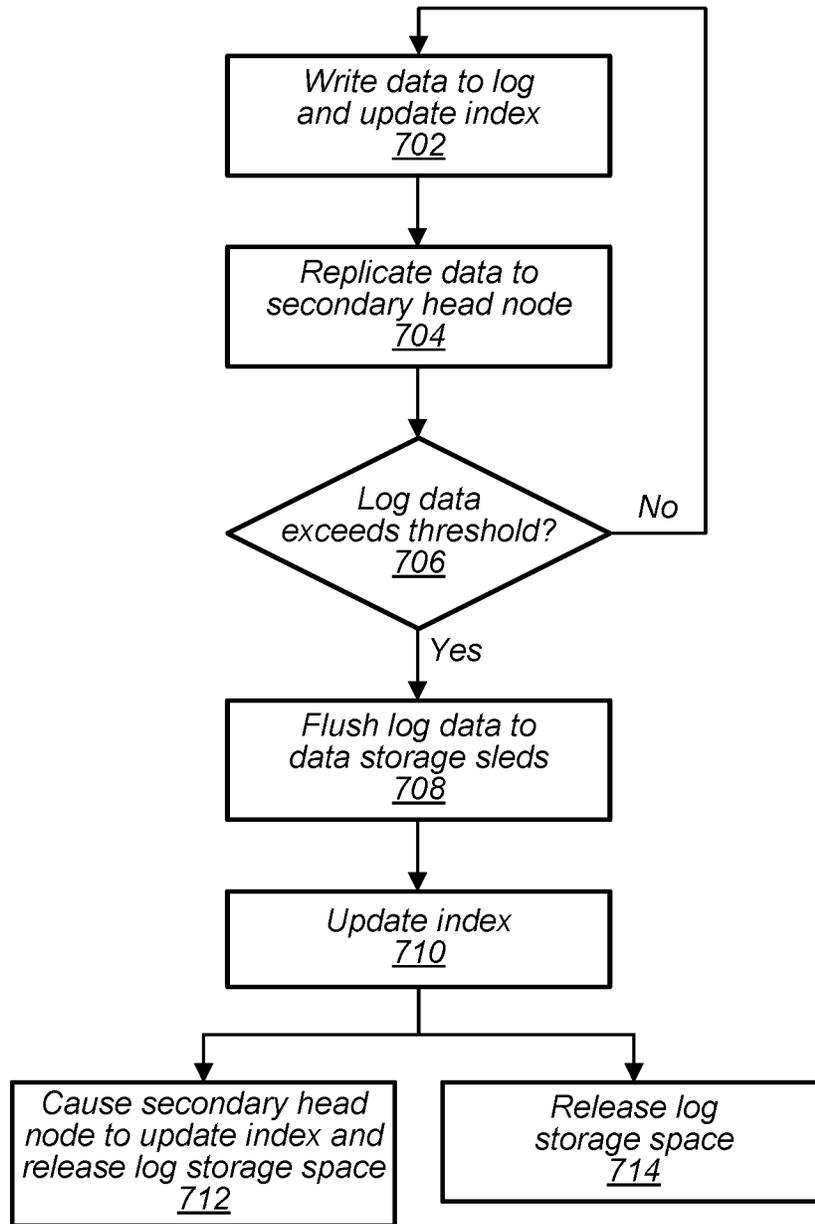


Fig. 7

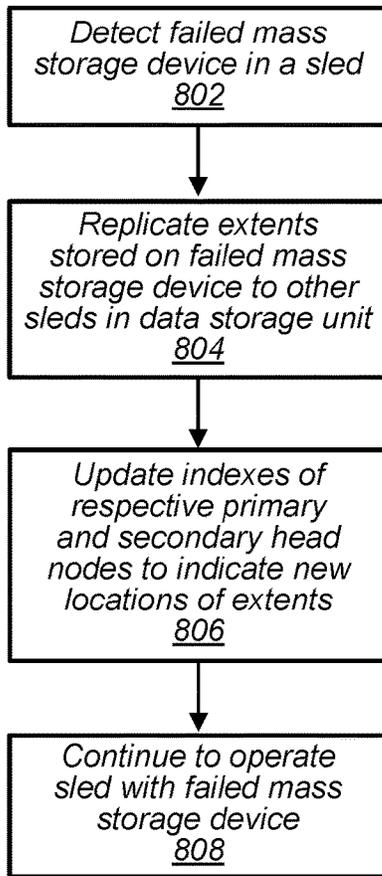


Fig. 8A

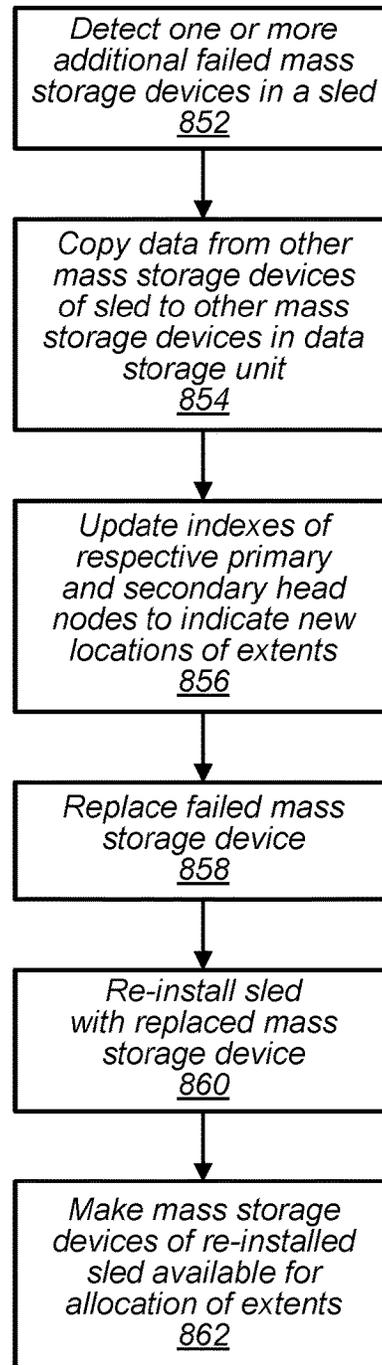


Fig. 8B

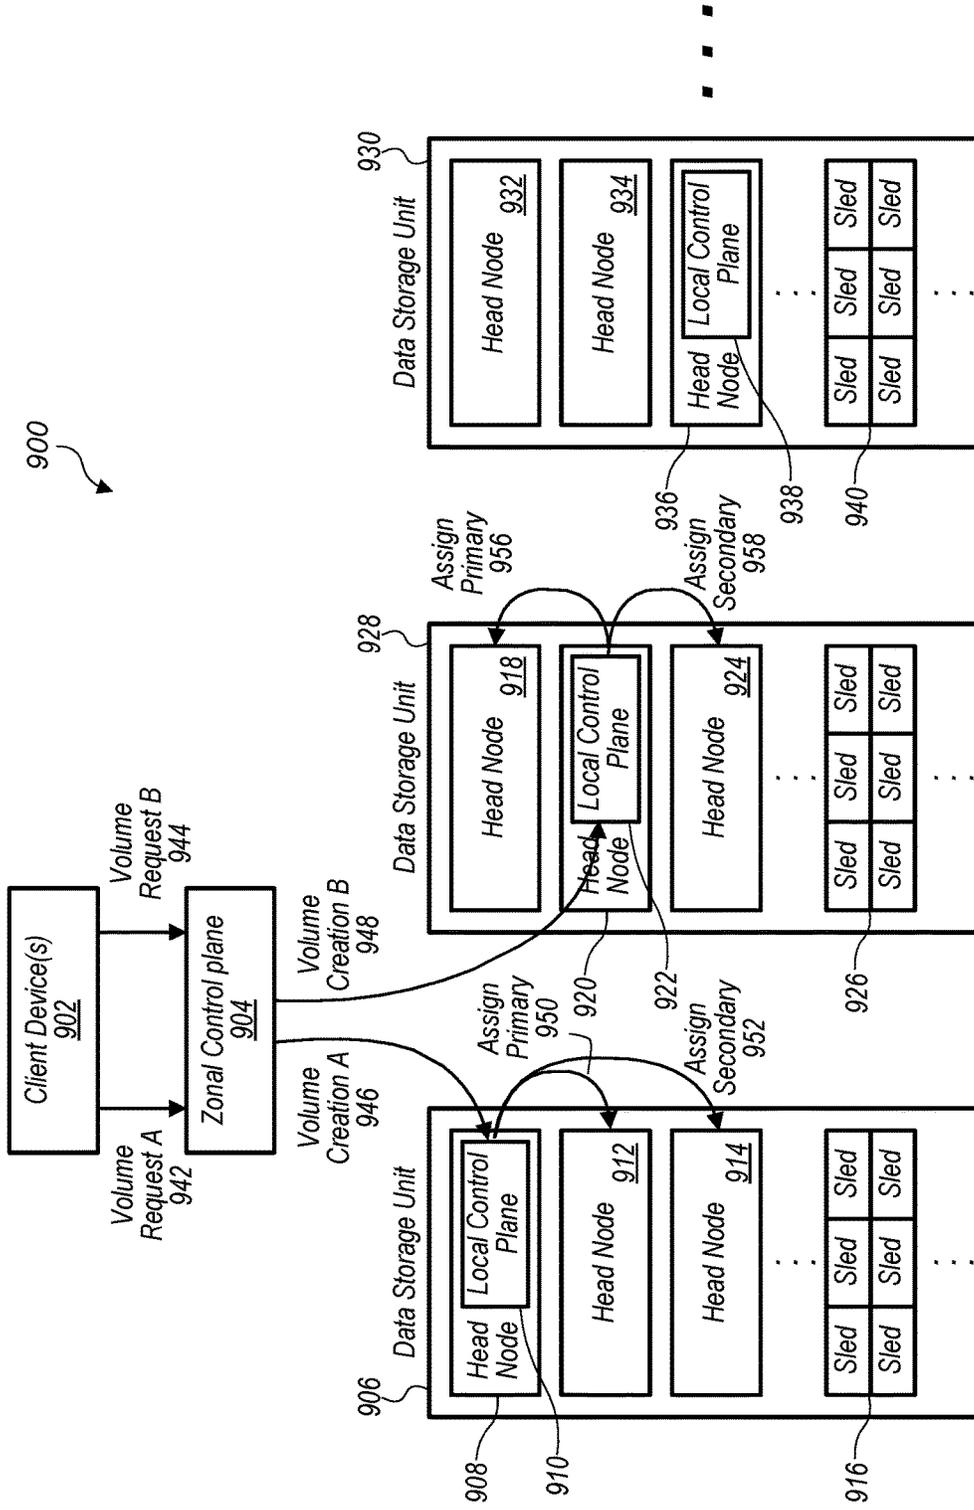


Fig. 9A

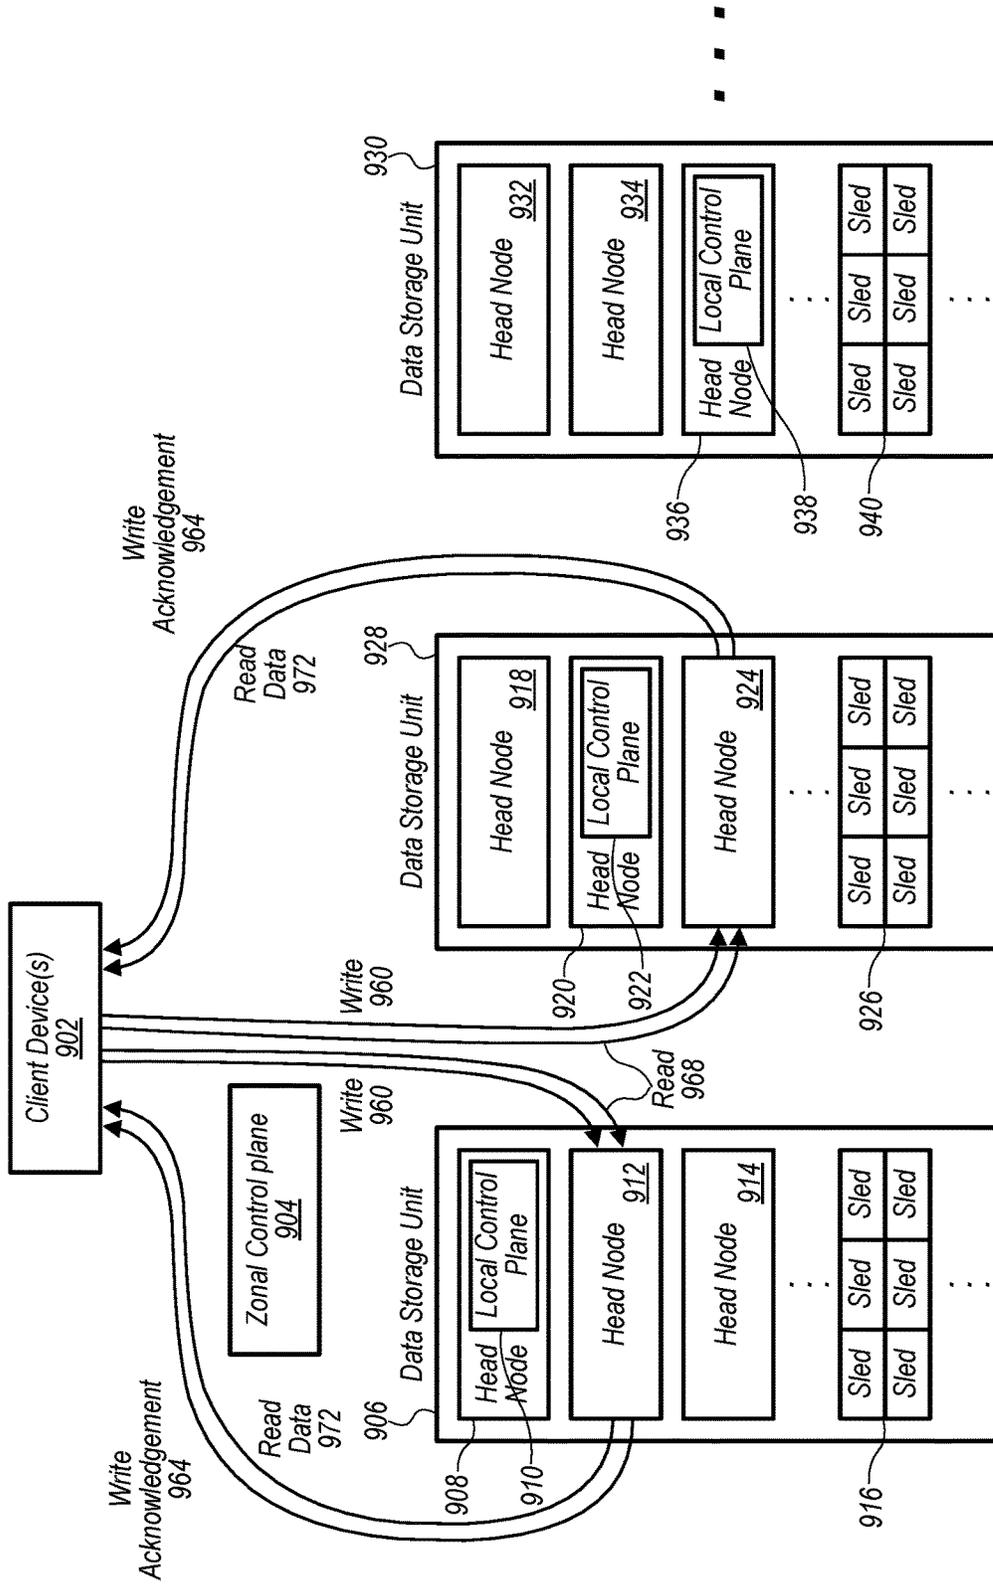


Fig. 9B

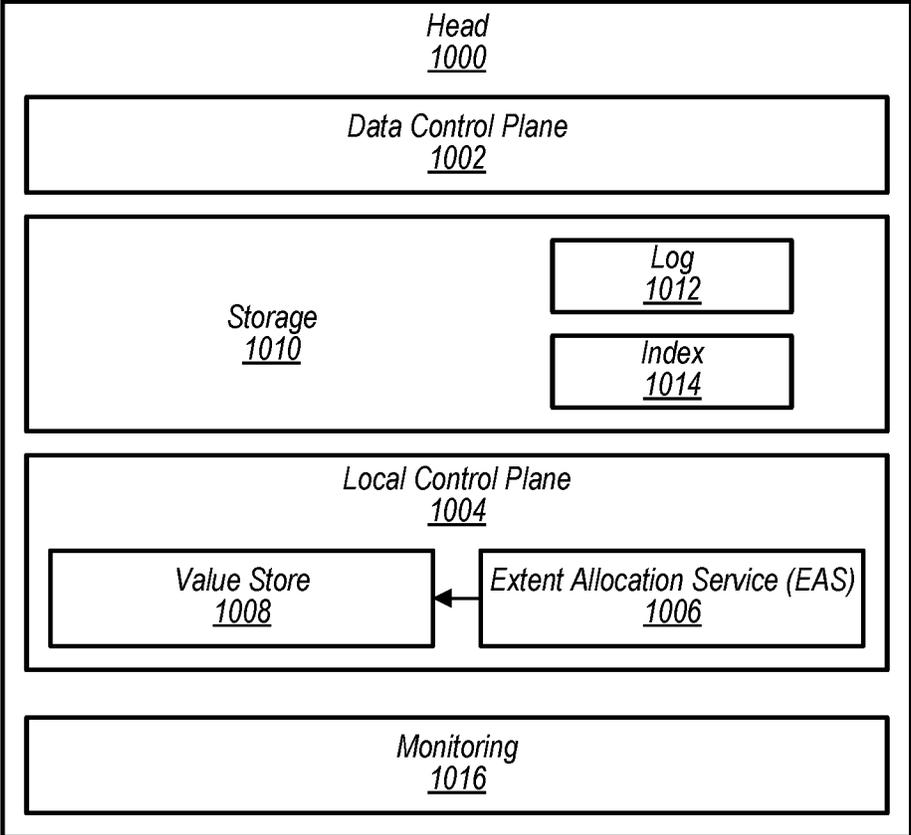


Fig. 10A

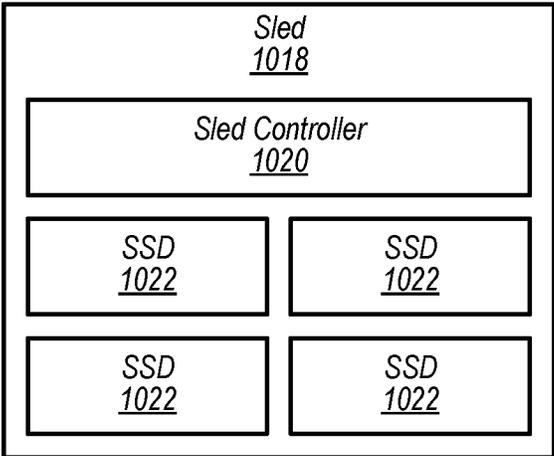


Fig. 10B

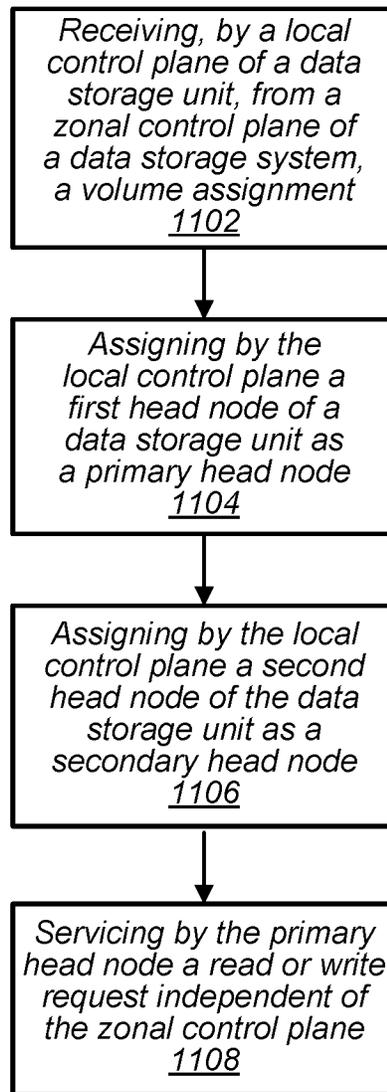


Fig. 11

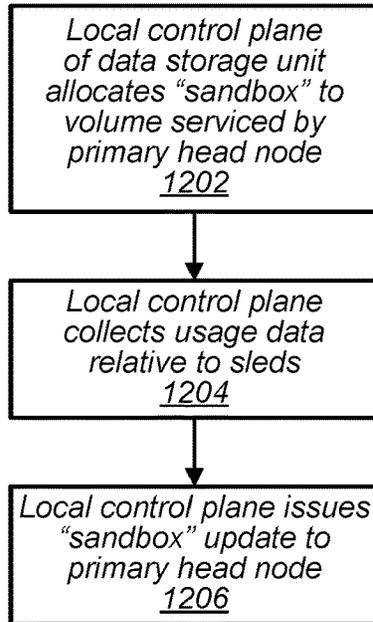


Fig. 12A

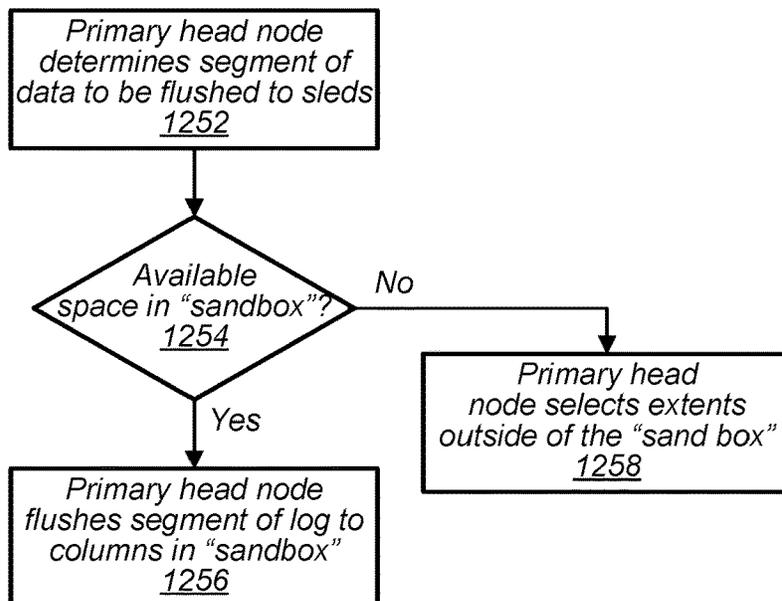


Fig. 12B

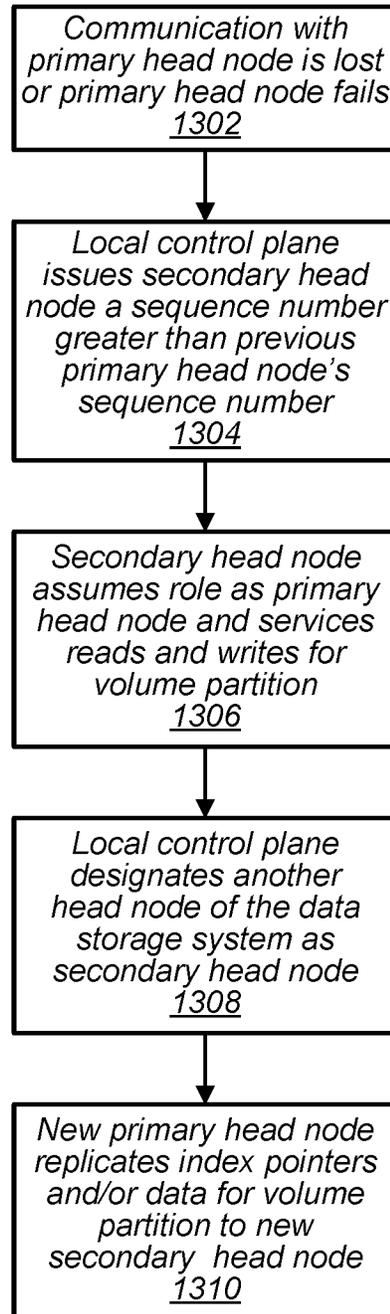


Fig. 13

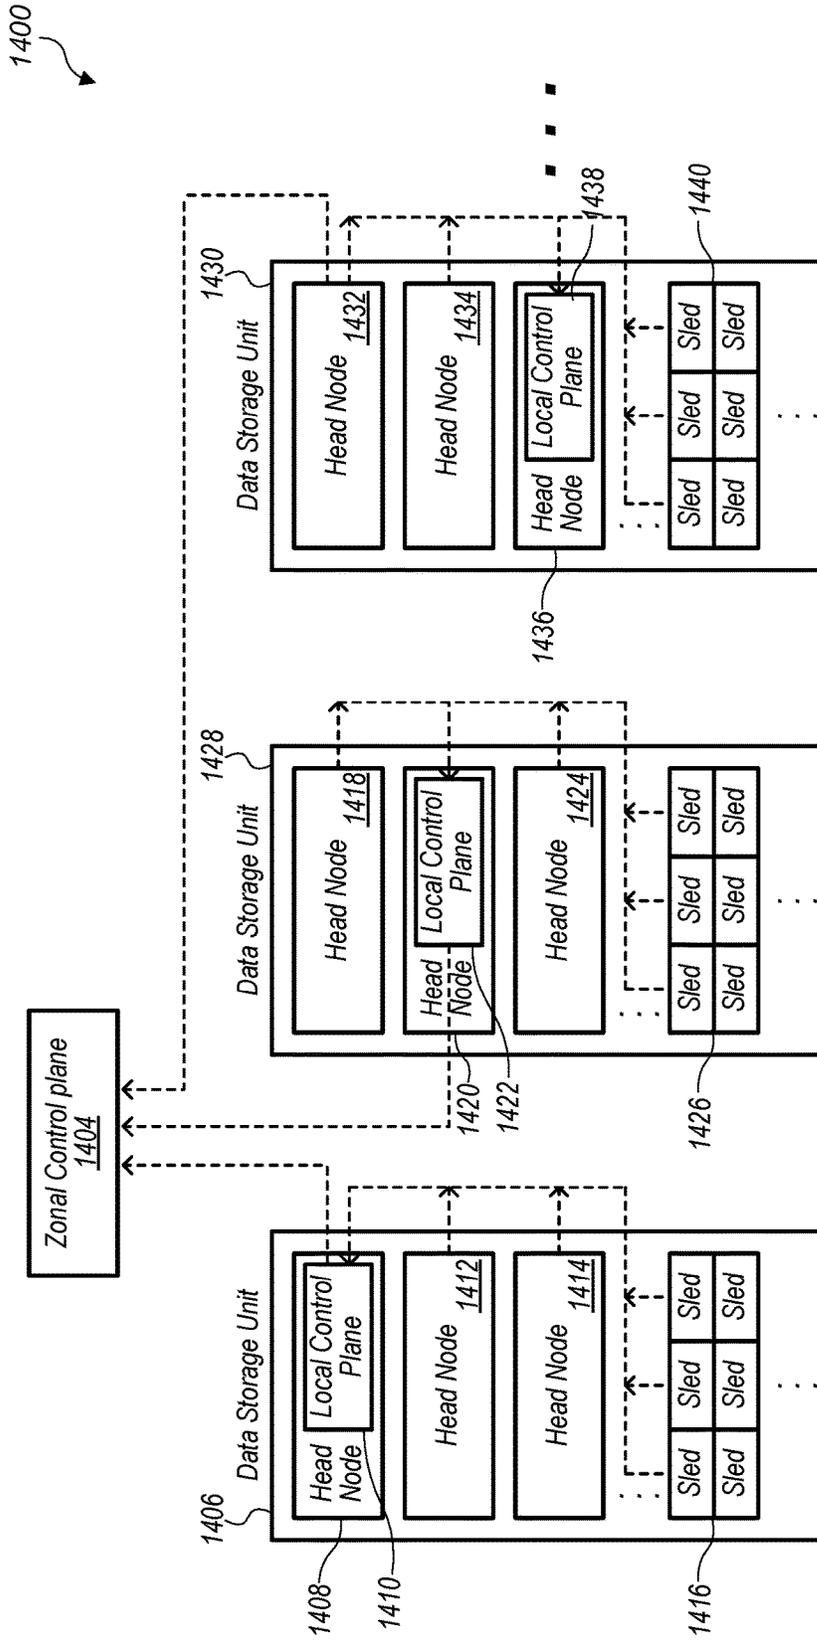


Fig. 14

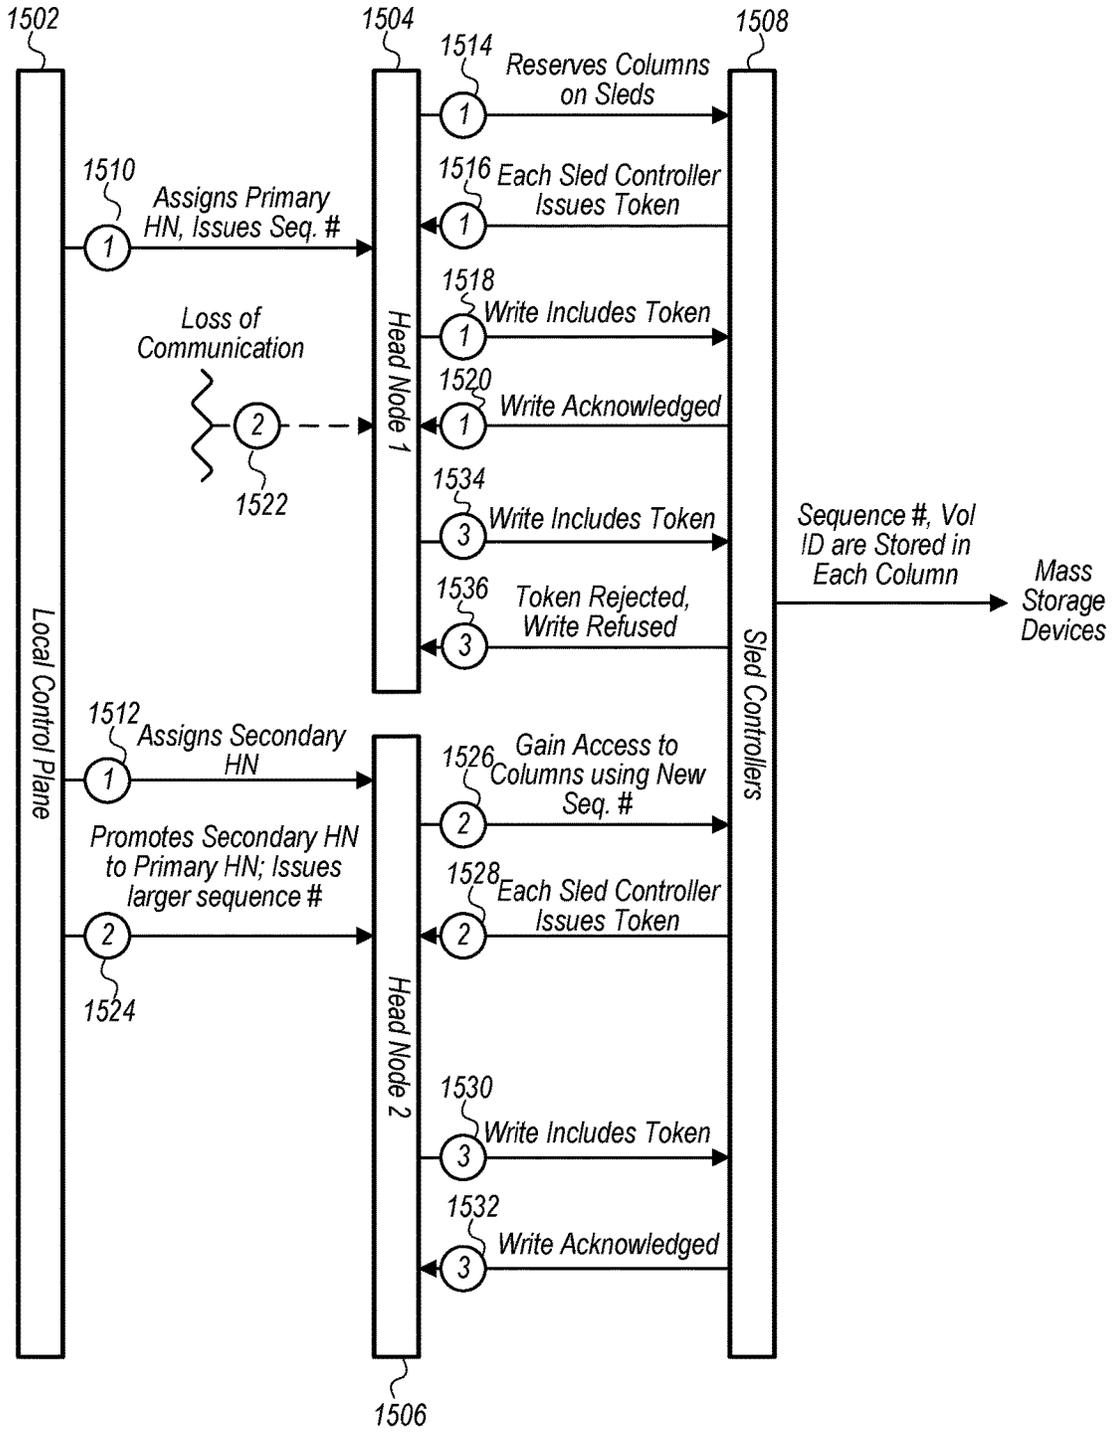
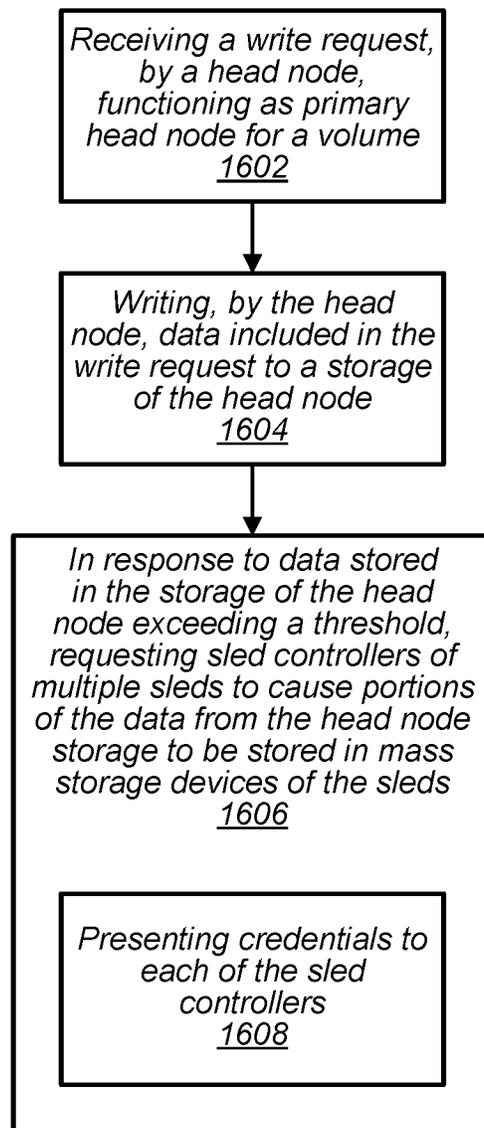


Fig. 15

*Fig. 16*

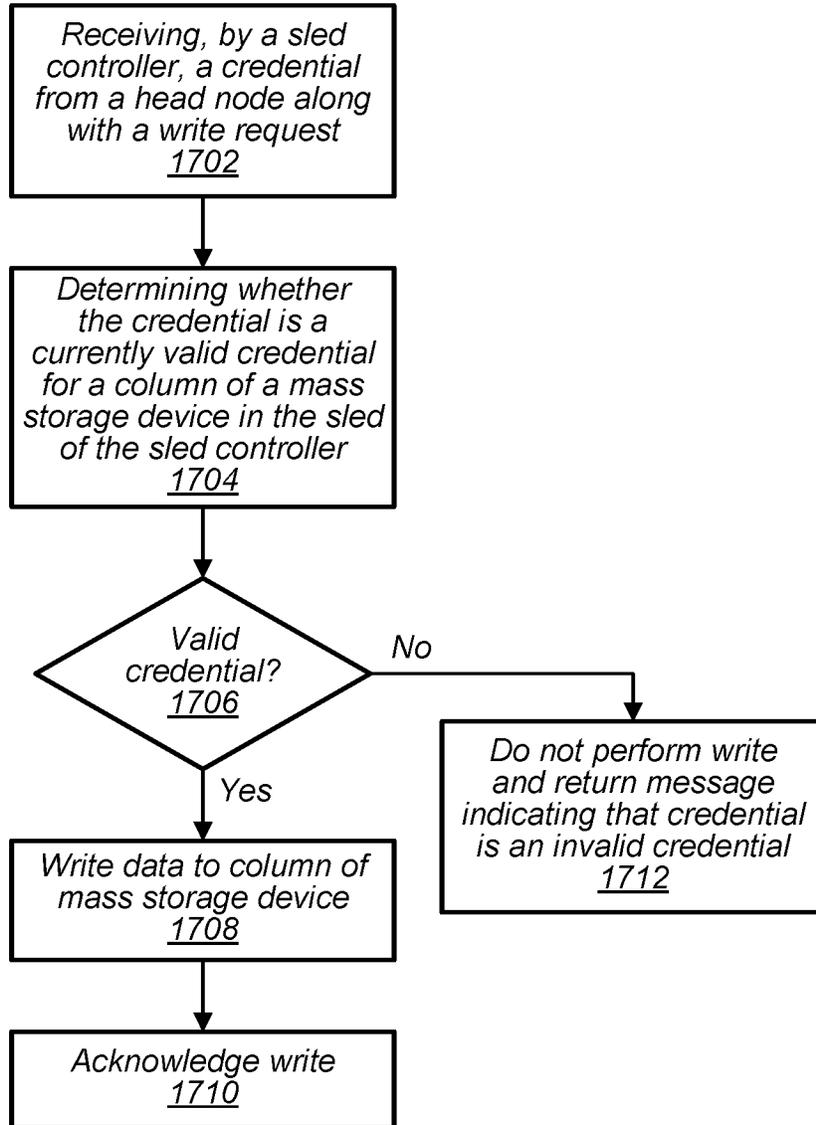


Fig. 17

Fig. 18A

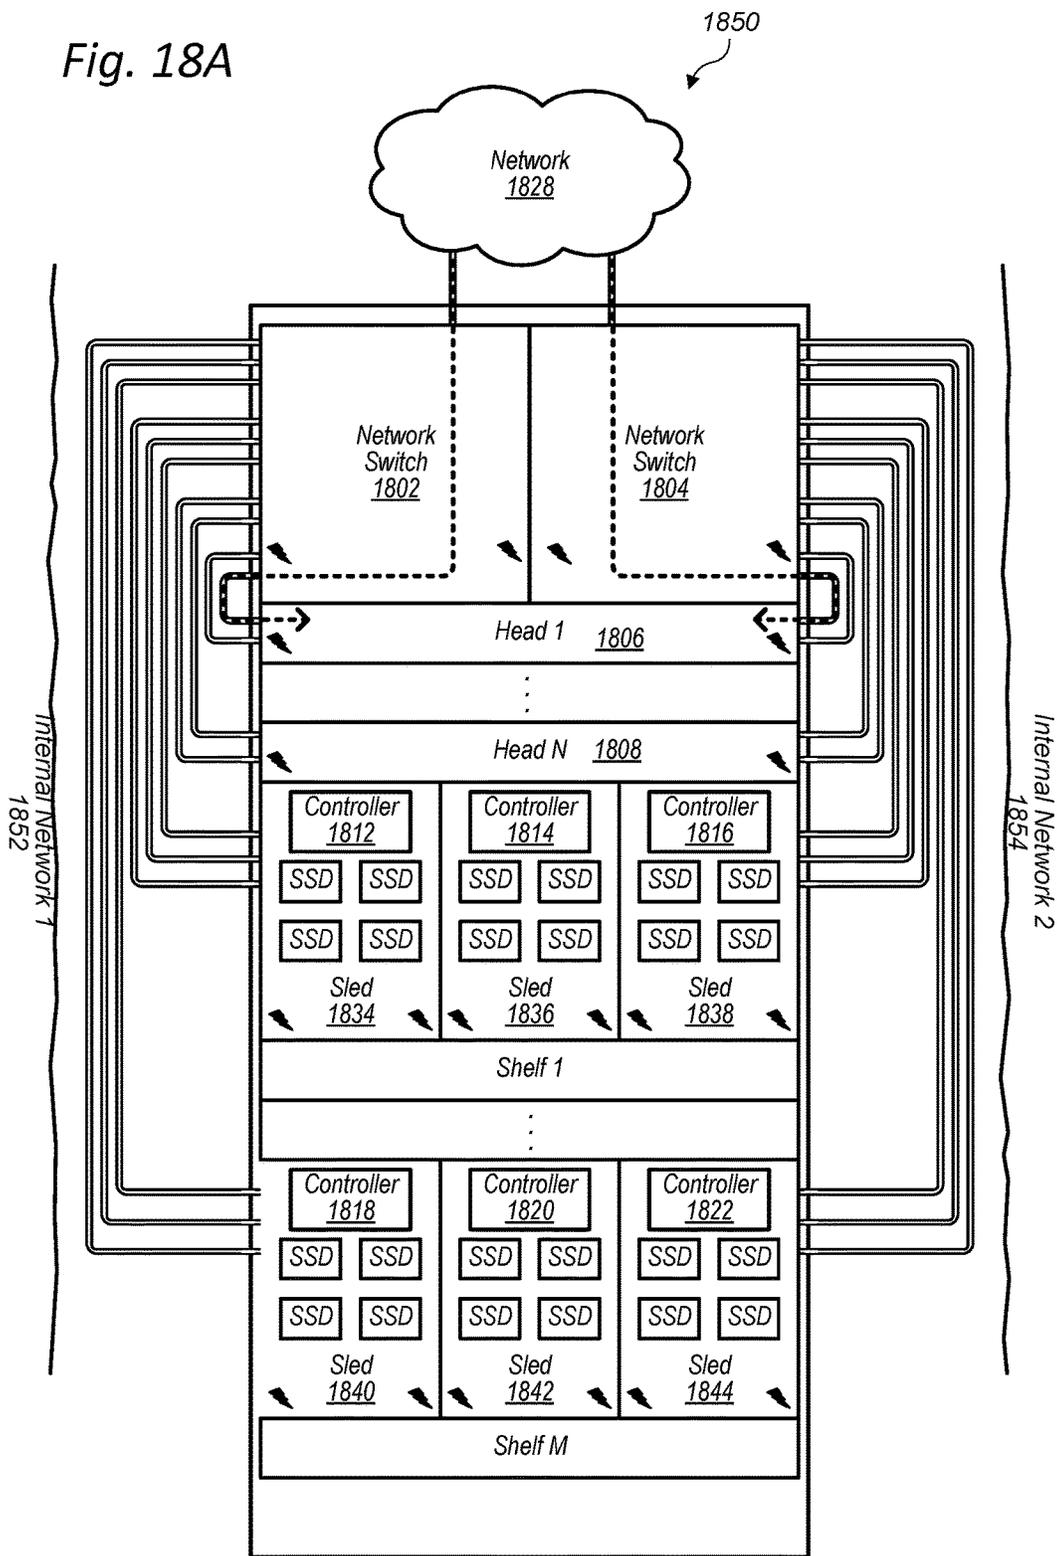


Fig. 18B

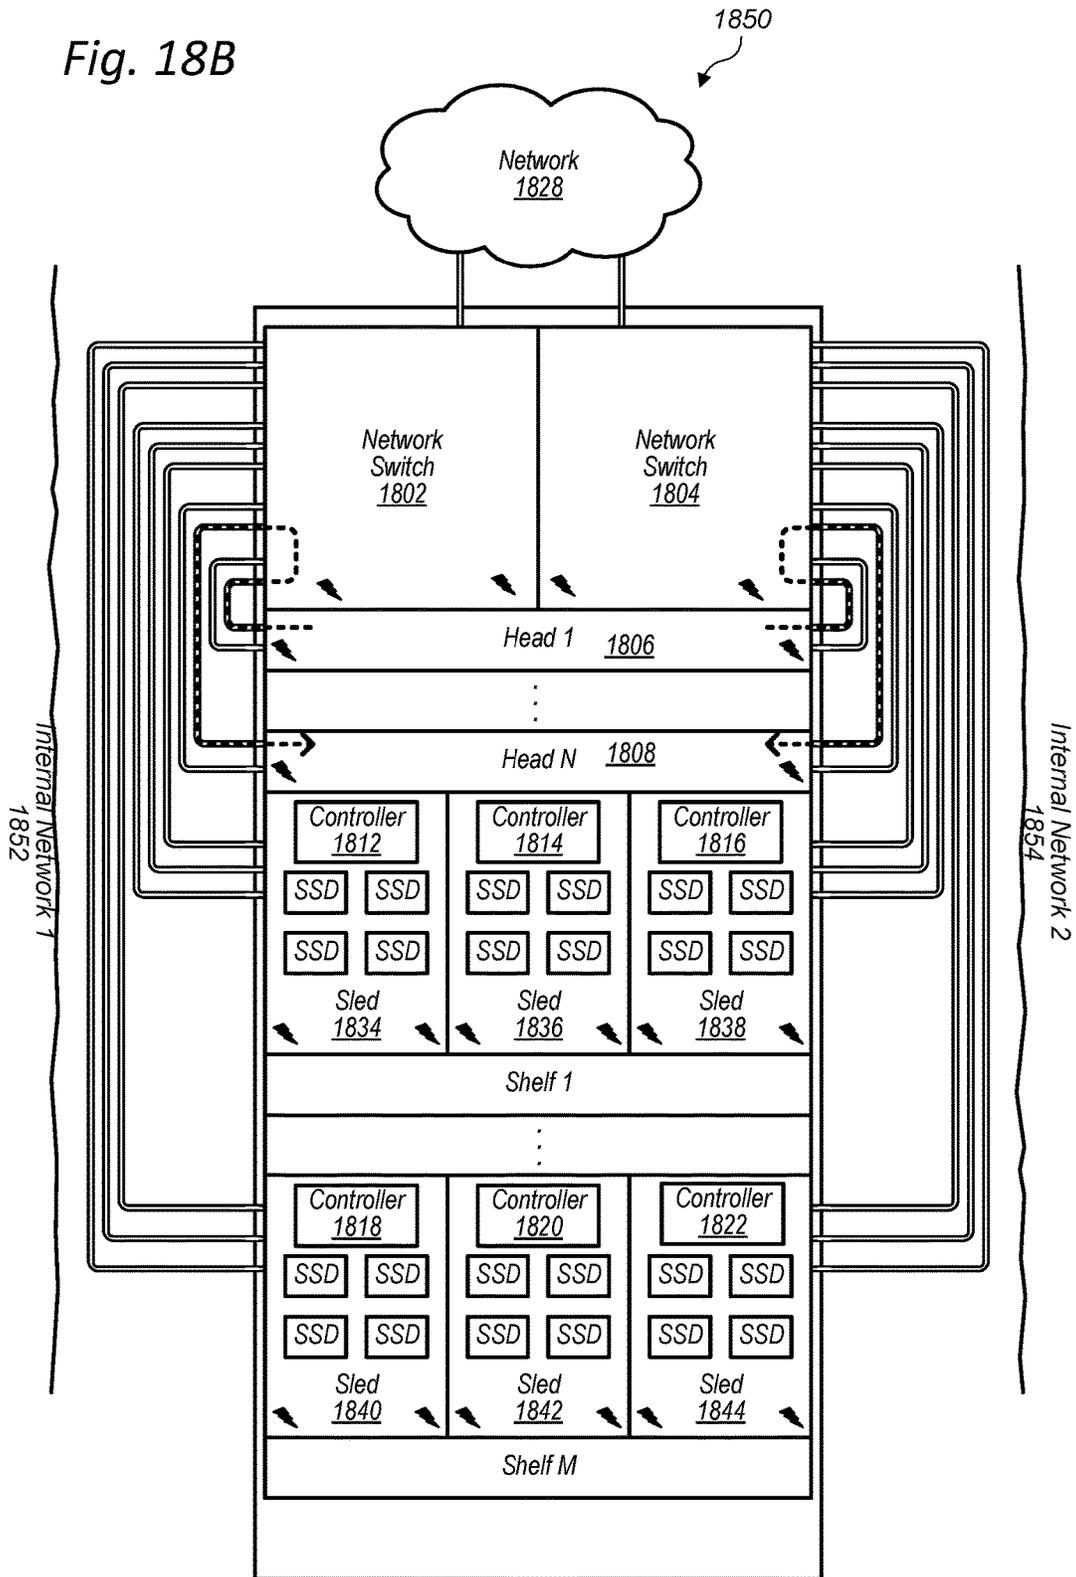


Fig. 18C

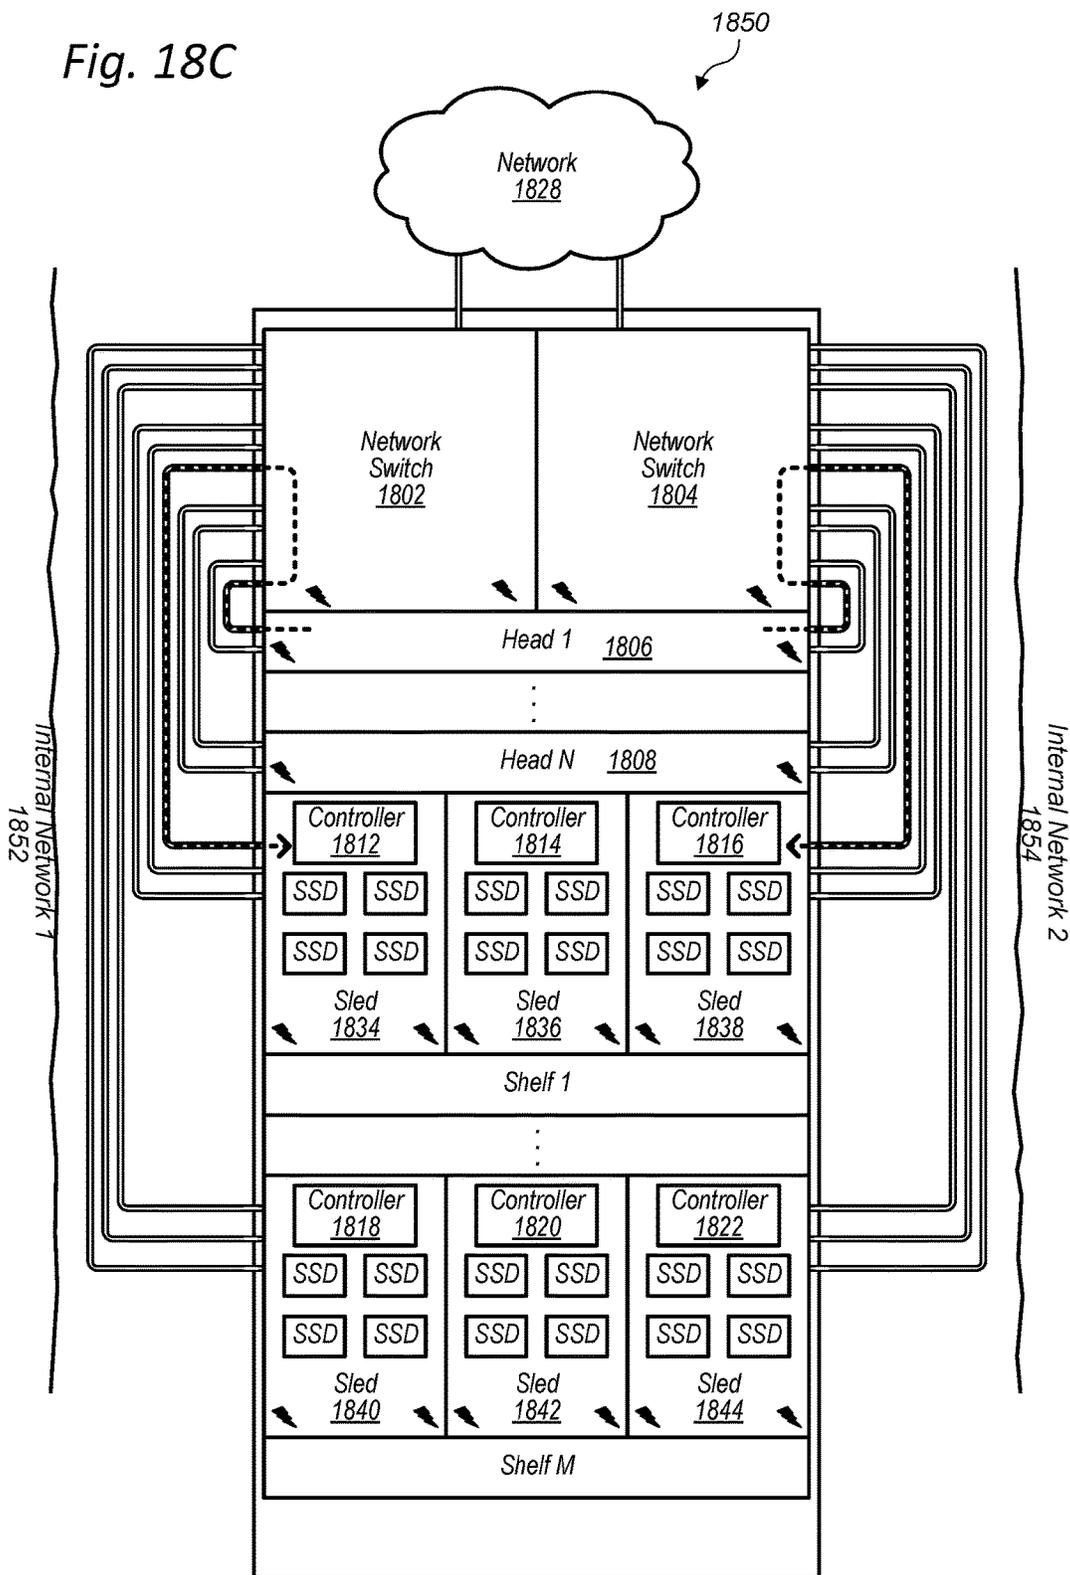
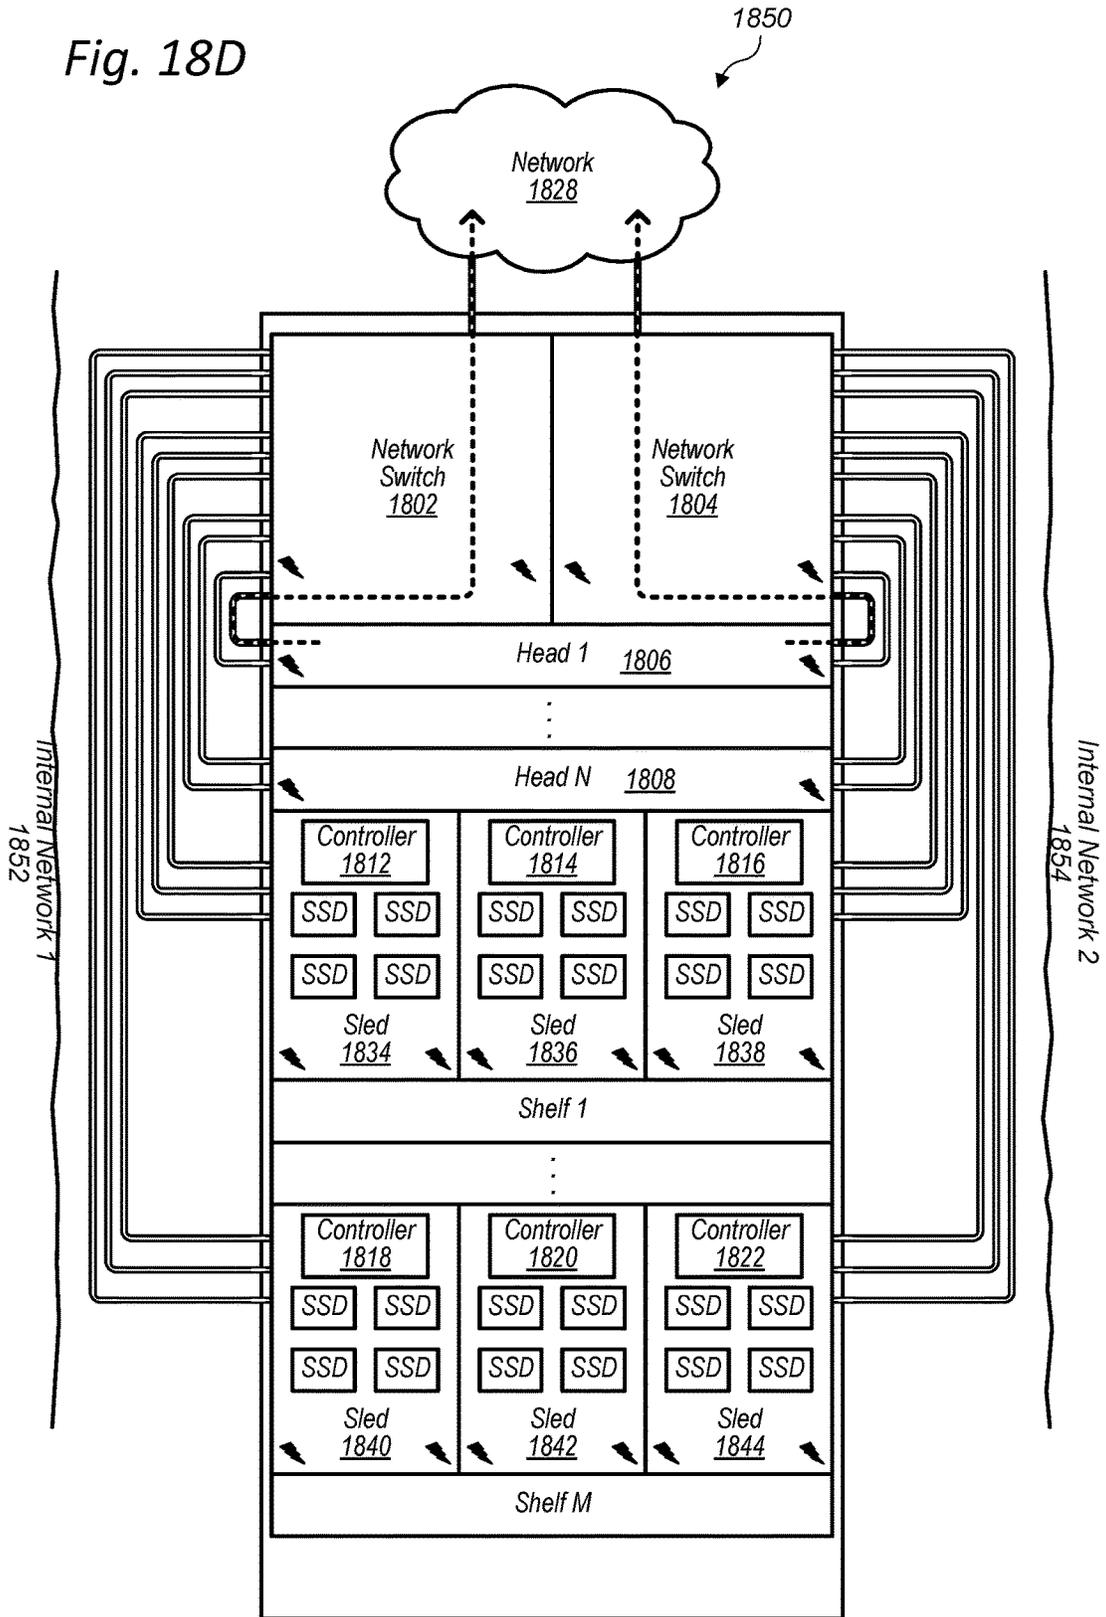


Fig. 18D



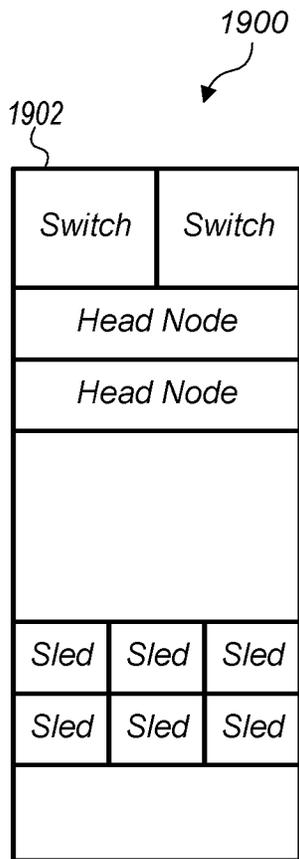


Fig. 19A

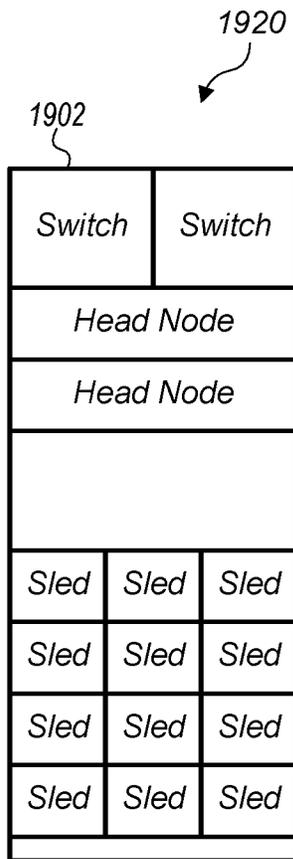


Fig. 19B

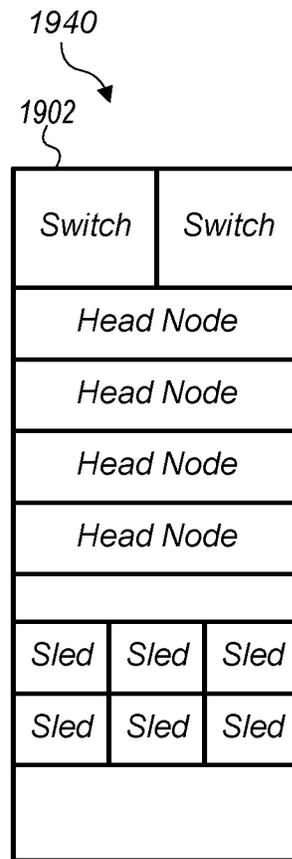


Fig. 19C

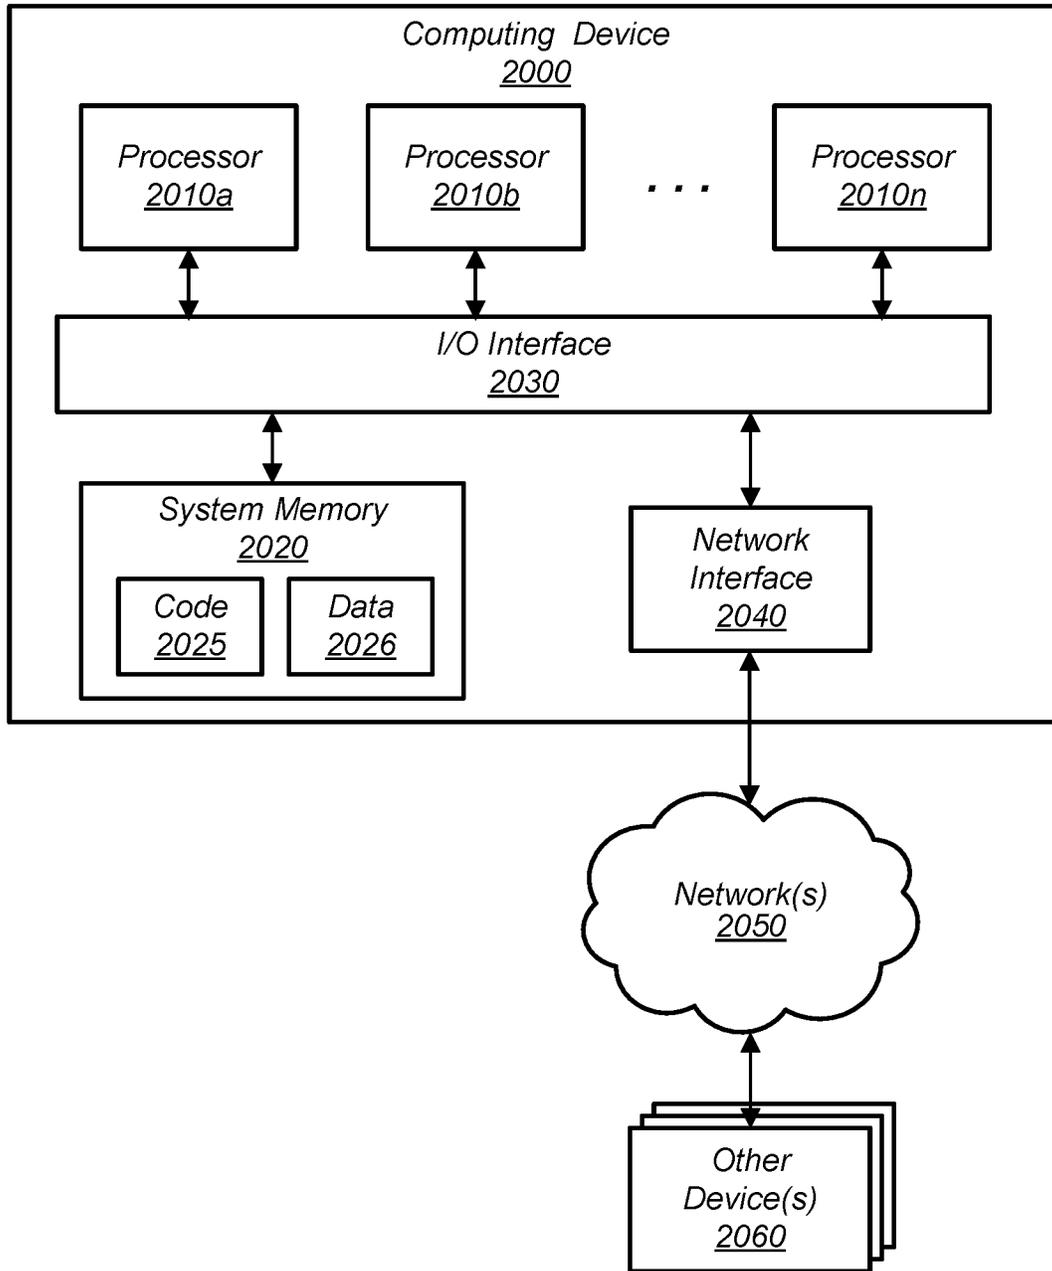


Fig. 20

DATA STORAGE SYSTEM WITH MULTIPLE DURABILITY LEVELS

BACKGROUND

[0001] The recent revolution in technologies for dynamically sharing virtualizations of hardware resources, software, and information storage across networks has increased the reliability, scalability, and cost efficiency of computing. More specifically, the ability to provide on demand virtual computing resources and storage through the advent of virtualization has enabled consumers of processing resources and storage to flexibly structure their computing and storage costs in response to immediately perceived computing and storage needs. Virtualization allows customers to purchase processor cycles and storage at the time of demand, rather than buying or leasing fixed hardware in provisioning cycles that are dictated by the delays and costs of manufacture and deployment of hardware. Rather than depending on the accuracy of predictions of future demand to determine the availability of computing and storage, users are able to purchase the use of computing and storage resources on a relatively instantaneous as-needed basis.

[0002] Virtualized computing environments are frequently supported by block-based storage. Such block-based storage provides a storage system that is able to interact with various computing virtualizations through a series of standardized storage calls that render the block-based storage functionally agnostic to the structural and functional details of the volumes that it supports and the operating systems executing on the virtualizations to which it provides storage availability.

[0003] Some block-based storage systems utilize a server node and multiple storage nodes that are serviced by the server node or dual server nodes that service multiple storage nodes. For example, a storage area network (SAN) may include such an architecture. However, in such systems, a failure of one or more of the server nodes may result in a large amount of storage capacity served by the server node(s) being rendered unusable or may result in significant decreases in the ability of the storage system to service read and write requests.

[0004] In order to increase durability of data, some block-based storage systems may store data across multiple devices in multiple locations. For example, a SAN may span multiple locations such as different facilities or different geographic locations. Such systems may utilize a common control plane to manage data in the multiple locations. However, in such systems, a failure of a component of the common control plane may impact a large quantity of storage capacity and render the large quantity of storage capacity unavailable. Also, such systems may require extensive networks to move data between the multiple locations and may also result in high latencies for data recovery due to data being located across the multiple locations.

BRIEF DESCRIPTION OF THE DRAWINGS

[0005] FIG. 1 illustrates a data storage unit comprising head nodes and data storage sleds, according to some embodiments.

[0006] FIG. 2 is a block diagram illustrating a provider network implementing multiple network-based services including a block-based storage service that includes data storage units, according to some embodiments.

[0007] FIG. 3 is a block diagram illustrating head nodes and data storage sleds of a data storage unit storing block storage data in response to a write request, according to some embodiments.

[0008] FIGS. 4A-4B are block diagrams illustrating a log storage and index of a head node storage, according to some embodiments.

[0009] FIG. 5 illustrates a partial view of a data storage unit that stores portions of a volume partition in multiple mass storage devices in multiple data storage sleds on multiple shelves of the data storage unit, according to some embodiments.

[0010] FIGS. 6A-B illustrate columns of mass storage devices storing different portions of a volume partition, according to some embodiments.

[0011] FIG. 7 is a high-level flowchart illustrating operations performed by a head node in response to a request to store data in a data storage unit, according to some embodiments.

[0012] FIG. 8A is a high-level flowchart illustrating operations performed by a head node in response to a failed mass storage device in a data storage sled of a data storage unit, according to some embodiments.

[0013] FIG. 8B is a high-level flowchart illustrating operations performed by a head node in response to a failed mass storage device in a data storage sled of a data storage unit, according to some embodiments.

[0014] FIG. 9A is a block diagram illustrating a process for creating a volume involving a zonal control plane, a local control plane, and head nodes of a data storage system, according to some embodiments.

[0015] FIG. 9B is a block diagram illustrating head nodes of a data storage unit servicing read and write requests independent of a zonal control plane of a data storage system, according to some embodiments.

[0016] FIG. 10A is a block diagram of a head node, according to some embodiments.

[0017] FIG. 10B is a block diagram of a data storage sled, according to some embodiments.

[0018] FIG. 11 is a high-level flowchart illustrating a process of creating a volume in a data storage system, according to some embodiments.

[0019] FIG. 12A is a high-level flowchart illustrating a local control plane of a data storage unit providing storage recommendations to a head node of the data storage unit for locations to store data in data storage sleds of the data storage unit for a volume serviced by the head node, according to some embodiments.

[0020] FIG. 12B is a high-level flowchart illustrating a head node of a data storage unit storing data in data storage sleds of the data storage unit, according to some embodiments.

[0021] FIG. 13 is a high-level flowchart illustrating head nodes of a data storage unit performing a fail over operation in response to a failure of or loss of communication with one of the head nodes of the data storage unit, according to some embodiments.

[0022] FIG. 14 is a block diagram illustrating performance and/or usage metrics being collected and accumulated in a data storage unit, according to some embodiments.

[0023] FIG. 15 illustrates interactions between a local control plane, head nodes, and data storage sleds of a data storage unit in relation to writing data to mass storage

devices of a data storage sled of a data storage unit, according to some embodiments.

[0024] FIG. 16 is a high-level flowchart of a head node of a data storage unit flushing data stored in a storage of the head node to a data storage sled of the data storage unit, according to some embodiments.

[0025] FIG. 17 is a high-level flowchart of a sled controller of a data storage sled processing a write request, according to some embodiments.

[0026] FIGS. 18A-D illustrate a data storage unit with redundant network paths within the data storage unit, according to some embodiments.

[0027] FIGS. 19A-C illustrate a data storage unit configured to allow scaling of storage capacity and processing capacity, according to some embodiments.

[0028] FIG. 20 is a block diagram illustrating an example computing system, according to some embodiments.

[0029] While embodiments are described herein by way of example for several embodiments and illustrative drawings, those skilled in the art will recognize that the embodiments are not limited to the embodiments or drawings described. It should be understood, that the drawings and detailed description thereto are not intended to limit embodiments to the particular form disclosed, but on the contrary, the intention is to cover all modifications, equivalents and alternatives falling within the spirit and scope as defined by the appended claims. The headings used herein are for organizational purposes only and are not meant to be used to limit the scope of the description or the claims. As used throughout this application, the word “may” is used in a permissive sense (i.e., meaning having the potential to), rather than the mandatory sense (i.e., meaning must). Similarly, the words “include”, “including”, and “includes” mean including, but not limited to.

DETAILED DESCRIPTION

[0030] According to one embodiment, a data storage system includes a rack, a plurality of head nodes mounted in the rack, and a plurality of data storage sleds mounted in the rack. For a partition of a volume to be stored in the data storage system, a particular one of the head nodes is designated as a primary head node for the volume partition and another one of the head nodes is designated as a secondary head node for the volume partition. In response to receiving a write request for the volume partition, the head node designated as the primary head node for the volume partition is configured to write data included with the write request to a storage of the head node designated as the primary head node and cause the data included with the write request to be replicated to the other head node designated as the secondary head node. Furthermore, the head node designated as the primary head node for the volume partition is further configured to cause respective parts of the data stored in the storage of the head node to be stored in a plurality of respective mass storage devices each in different ones of the plurality of data storage sleds of the data storage system. For example, a data storage system may store data in a storage of a primary head node and replicate the data to a storage of a secondary head node. Then, after a certain amount of time has passed, a certain amount of data has been written for the volume partition, or in response to another trigger, the head node may cause the data stored in the storage of the head node to be stored in multiple mass storage devices of different ones of the data storage sleds of the data storage

system. For example, data may be stored in mass storage devices of different data storage sleds of a data storage system in a RAID array and may be erasure encoded across the multiple mass storage devices. Such a system may provide varying latencies for accessing stored data and different durabilities of the stored data based on whether the data is stored in storages of the primary and secondary head nodes or stored in multiple mass storage devices of multiple data storage sleds of the data storage system. For example, data stored in a storage of a primary head node may be accessed with lower latencies than data stored across multiple data storage sleds of a data storage system. However, data stored across multiple data storage sleds of a data storage system may have higher durability than data replicated between storages of a primary and secondary head node. Thus, a data storage system may provide low latencies for recently or frequently accessed data while providing high durability for long term storage of data or for data less frequently accessed. In some embodiments, durability of data stored and replicated in head nodes may be adjusted by varying a number of head nodes that replicate the data. Also, durability of data stored in mass storage devices of data storage sleds of a data storage system may be adjusted by varying a RAID scheme or data encoding procedure used to store the data amongst other techniques to increase data durability.

[0031] According to one embodiment, a data storage system includes a head node of a data storage system, wherein a plurality of data storage sleds of the data storage system are also included in the data storage system. The head node, when acting as a primary head node of the data storage system for the volume partition and in response to receiving a write request for a volume partition, is configured to write data included with the write request to a storage of the head node and cause the data included with the write request to be replicated to another head node of the data storage system designated as a secondary head node for the volume partition. The head node is further configured to cause respective parts of the data stored in the storage of the head node to be stored in a plurality of respective mass storage devices each in different ones of the plurality of data storage sleds mounted in the rack of the data storage system.

[0032] According to one embodiment, a non-transitory computer readable medium stores program instructions for implementing a head node of a data storage system, wherein the program instructions when executed by a processor cause the system to, in response to receiving a write request for a volume partition, write data included with the write request to a storage of a head node designated as primary head node for the volume partition and cause the data included with the write request to be replicated to another head node of the data storage system designated as a secondary head node for the volume partition. The program instructions when executed by the processor further cause respective parts of the data stored in the storage of the head node designated as the primary head node to be stored in a plurality of respective mass storage devices each in different ones of a plurality of data storage sleds mounted in a rack of the data storage system.

[0033] According to one embodiment, a data storage system includes a plurality of data storage units, for example that are each hosted on a rack, a plurality of head nodes mounted in the rack, and a plurality of data storage sleds. The data storage system also includes one or more comput-

ing devices external to the plurality of data storage units configured to implement a zonal control plane for partially controlling storage operations related to the plurality of data storage units. In response to a volume creation request, the zonal control plane is configured to assign a particular one of the data storage units to service a volume requested by the volume creation request. Also, for each respective data storage unit, at least one of the head nodes of the respective data storage unit is configured to implement a local control plane for the respective data storage unit, wherein the plurality of head nodes are configured to service read requests and write requests directed to one or more volumes stored in the respective data storage unit independent of the local control plane and the zonal control plane. For example, a data storage system may include data storage units that are configured to service read and write requests without the read and write requests being routed through a local control plane or a zonal control plane of the data storage system. Also, the data storage units of the data storage system may continue to service read and write requests from client devices regardless of whether communication with a local control plane or a zonal control plane of the data storage system is available or lost.

[0034] According to one embodiment, a data storage system includes a data storage unit comprising a plurality of head nodes and a plurality of data storage sleds. At least one of the head nodes of the data storage unit implements a local control plane for the data storage unit, wherein, in response to a volume creation request, the local control plane is configured to receive an assignment from a zonal control plane to service a volume requested by the volume creation request, wherein the zonal control plane at least partially controls storage operations related to the data storage unit and one or more additional data storage units. Furthermore, the plurality of head nodes of the data storage unit are configured to service read requests and write requests directed to one or more volumes serviced by the data storage unit independent of the zonal control plane. For example, once a volume is created on a particular data storage unit in response to a zonal control plane receiving a volume creation request, the data storage unit may service read and write requests directed to the volume independent of the zonal control plane.

[0035] According to one embodiment, a method includes receiving, from a zonal control plane by a local control plane of a data storage unit, an assignment of a volume to be serviced by the data storage unit, wherein the data storage unit comprises a plurality of head nodes and a plurality of data storage sleds, wherein at least one of the head nodes implements the local control plane of the data storage unit. The method further includes assigning, by the local control plane, a particular one of the head nodes of the data storage unit to function as a primary head node for the volume; assigning, by the local control plane, another particular one of the head nodes of the data storage unit to function as a secondary head node from the volume; and in response to a read or write request for the volume, servicing, by the primary head node of the data storage unit, the read or write request independent of the zonal control plane.

[0036] According to one embodiment a data storage system includes a plurality of head nodes and a plurality of data storage sleds. Each of the data storage sleds includes multiple mass storage devices and a sled controller for the plurality of mass storage devices mounted in the data storage

sled. Respective ones of the head nodes are configured to obtain credentials for accessing particular portions of the mass storage devices of respective ones of the plurality of data storage sleds. For example, a head node may receive a credential from a local control plane implemented on one or more of the head nodes of the data storage unit or may receive credentials from a zonal control plane implemented on one or more computing devices external to the data storage unit. Each of the respective sled controllers, in response to a request from a particular head node to write data on a particular portion of a particular mass storage device in a particular data storage sled that includes the respective sled controller, is configured to determine whether a credential included with the write request from the particular head node is a valid credential for accessing the particular portion of the particular mass storage device. In response to determining the credential is a valid credential for the particular portion of the particular mass storage device, the respective sled controller is configured to cause the requested write to be performed on the particular portion of the particular mass storage device. Also, in response to determining the credential is an invalid credential for the particular portion of the particular mass storage device, the respective sled controller is configured to decline to perform the requested write and return a message to the particular head node indicating the credential for accessing the particular portion of the particular mass storage device is an invalid credential. For example, if a credential for writing to a particular portion of a mass storage device is issued to a head node functioning as a primary head node for a volume and another head node of the data storage unit attempts to write to the particular portion of the mass storage device without a credential or with an inferior credential that is inferior to the credential held by the primary head node, the sled controller of the data storage sled may enforce the fencing off of the particular portion of the mass storage device for the head node functioning as the primary head node for the volume by refusing to perform the write requested by the other head node of the data storage unit. Also, in some embodiments, a head node functioning as a primary head node may determine that it has been superseded as primary head node by another head node of a data storage unit in response to a write request being denied by a sled controller. Such a scheme may prevent corruption of data caused by a head node attempting to write to a particular portion of a mass storage device after another head node of a data storage unit has taken over as primary head node and assumed exclusive responsibility for writing new data to the particular portion of the mass storage device.

[0037] According to one embodiment, a method includes receiving, by a head node of a data storage system, a write request from a client of the data storage system; writing, by the head node, data included with the write request to a data storage of the head node; and requesting, by the head node, to write the data included with the write request to a plurality of mass storage devices in a plurality of data storage sleds of the data storage system. Requesting to write the data to the data storage sleds includes presenting respective credentials to respective sled controllers of each of the plurality of data storage sleds, wherein the respective sled controllers cause the data to be written to the respective mass storage devices of the respective data storage sleds in response to

determining the respective credentials are valid credentials for accessing respective portions of the respective mass storage devices.

[0038] According to one embodiment, a method includes determining, by a sled controller of a data storage system, whether a credential included with a write request from a particular head node of the data storage system is a valid credential for accessing a particular portion of a particular mass storage device included in a sled with the sled controller. The method further includes in response to determining the credential is a valid credential for the particular portion of the particular mass storage device, causing the requested write to be performed on the particular portion of the particular mass storage device. The method also includes determining, by a sled controller of a data storage system, whether another credential included with another write request from another particular head node of the data storage system is a valid credential for accessing the particular portion of the particular mass storage device included in the sled with the sled controller; and in response to determining the other credential is an invalid credential for the particular portion of the particular mass storage device, declining to perform the requested write and returning a message to the other particular head node indicating the other credential for accessing the particular portion of the particular mass storage device is an invalid credential.

[0039] According to one embodiment, a data storage system comprises a plurality of head nodes, for example mounted on a rack, a plurality of data storage sleds, and at least two networking devices. The at least two networking devices are configured to implement at least two redundant networks within the data storage system, wherein to implement the at least two redundant networks each respective head node is coupled to each of the plurality of data storage sleds via a first one of the at least two networking devices, each respective head node is also coupled to each of the plurality of data storage sleds via a second one of the at least two networking devices, and each respective head node is assigned at least two unique network addresses for communicating with the plurality of data storage sleds. For example, a particular head node of a data storage unit may be configured to communicate with external devices via a first path through the first networking device and using a first address, such as a first IP address and also communicate with the external device via a redundant network path through the second networking device and using a second address, such as a second IP address. Also, a head node may be configured to communicate with mass storage devices in separate ones of the data storage sleds mounted in the rack via a first path through the first networking device and through a second path through the second networking device. In some embodiments, a data storage unit may be configured such that only a single network hop is required for a head node to retrieve data stored in data storage sleds of the data storage unit.

[0040] According to one embodiment, a data storage system includes a head node comprising at least three network interfaces, wherein at least two of the network interfaces are configured to implement at least two redundant networks within the data storage system and at least one network interface of the at least three network interfaces is configured to enable communications between the data storage system and external clients. To implement the at least two redundant networks the at least two network interfaces of the

head node are configured to couple to each of a plurality of data storage sleds via a first networking device and couple to each of the plurality of data storage sleds via a second networking device. Also, the head node is assigned at least two unique network addresses for communicating with the plurality of data storage sleds.

[0041] According to one embodiment, a data storage system includes a data storage sled comprising at least two network interfaces configured to implement at least two redundant networks within the data storage system, wherein to implement the at least two redundant networks, the at least two network interfaces of the data storage sled are configured to couple to a particular head node via a first a networking device and couple to the same particular head node via a second networking device. Also, the data storage sled is assigned at least two unique network addresses for communicating with the particular head node or one or more additional head nodes of the data storage system.

[0042] Some data storage systems, such as storage area networks (SAN) may allow a server or a pair of servers to access a shared set of storage resources. However, such systems may be susceptible to significant losses in performance due to a server failure. Also, in such systems, data may be durably stored in storage devices of the SAN network, but not durably stored in the servers accessing the SAN network.

[0043] In order to provide high durability data storage and low latencies for accessing data, a data storage unit may store data in local storages of head nodes that function as servers for the data storage system, replicate the data to another head node of the data storage unit, and also store the data across multiple mass storage devices in multiple data storage sleds of the data storage unit. Thus, a data storage system that includes a data storage unit may provide low latency input/output operations for data stored in a storage of a head node, while still providing data durability due to the data being replicated to another head node. Furthermore, the data storage system may provide even higher durability for the data once the data is stored in multiple mass storage devices in different data storage sleds of the data storage unit. Thus, a data storage system may provide varying levels of data durability and input/output operation latency depending on whether the data is stored in a storage of a head node and replicated to another head node or whether the data is stored in multiple mass storage devices in different data storage sleds of the data storage system.

[0044] In some embodiments, data may be initially stored in a storage of a head node and replicated to a storage of another head node, and may be asynchronously moved to multiple mass storage devices in different data storage sleds that form a RAID array (random array of independent disks) to store the data. In some embodiments, recently stored data or frequently accessed data may remain in a head node storage to allow for low latency access to the data. The data may then be moved to mass storage devices in data storage sleds of a data storage unit of the data storage system after a certain amount of time has elapsed since the data was last accessed or stored. Moving the data to the mass storage devices may increase the durability of the data as compared to being stored in a storage of a primary head node and being replicated to a storage of a secondary head node. Thus a data storage system may provide different levels of durability and latency based on a staleness or a frequency of access to data stored in the data storage system. In some embodiments,

other criteria may be used to determine when data stored in a storage of a head node is to be moved to mass storage devices of data storage sleds of a data storage unit. For example, data may be collected in a log of a head node and upon an amount of data being stored in the log exceeding a threshold amount, the data may be moved to mass storage devices of data storage sleds of a data storage unit of the data storage system.

[0045] In some embodiments, a data storage unit of a data storage system may include multiple head nodes, multiple data storage sleds, and at least two networking devices. The data storage unit may further include connectors for coupling the data storage unit with at least two separate power sources. The data storage unit may also include at least two power distribution systems within the data storage unit to provide redundant power to the head nodes, the data storage sleds, and the networking devices of the data storage unit. Furthermore, the at least two networking devices of the data storage unit may implement at least two redundant networks within the data storage unit that enable communications between the head nodes of the data storage unit and the data storage sleds of the data storage unit. Furthermore, the at least two networking devices of the data storage unit may implement at least two redundant networks within the data storage unit that enable communications between the head nodes of the data storage unit and external clients of the data storage unit. In some embodiments, a data storage unit that include redundant networks and redundant power may provide high reliability and data durability for data storage and access while storing data locally within devices mounted within a single rack.

[0046] In some embodiments, a data storage unit of a data storage system may include multiple head nodes that are assigned network addresses that are routable from devices external to the data storage unit. Thus, external clients may communicate directly with head nodes of a data storage unit without the communications being routed through a control plane of the data storage system that is external to the data storage unit, such as a zonal control plane. Also, a data storage system that includes multiple data storage units may implement a zonal control plane that assigns volumes or volume partitions to particular ones of the data storage units of the data storage system. Also, a zonal control plane may coordinate operations between data storage units, such as rebalancing loads by moving volumes between data storage units. However, a data storage unit may also implement a local control plane configured to perform fail over operations for head nodes and mass storage devices of data storage sleds of the data storage unit. Because head nodes of a data storage unit may communicate directly with client devices and because a local control plane may manage fail over operations within a data storage unit, the data storage unit may operate autonomously without relying on a zonal control plane once a volume has been created on the data storage unit.

[0047] In some embodiments, in order to prevent corruption of data stored in mass storage devices of a data storage system, a data control plane may be at least partially implemented on a sled controller of a data storage sled of the data storage system. The data storage sled may include multiple mass storage devices serviced by the sled controller. Also, portions of respective mass storage devices of a particular data storage sled may be reserved for a particular volume serviced by a particular head node functioning as a

primary head node for the particular volume. In order to reserve the portions for the particular volume or a volume partition of the particular volume, a sled controller of a data storage sled may provide a token to a head node requesting to reserve the portions. Once the portions are reserved for the particular volume by the head node acting as the primary head node, the head node while acting as a primary head node for the particular volume, may provide the token to the sled controller along with a write request when writing new data to the portions. The sled controller may verify the token and determine the head node is authorized to write to the portions. Also, the sled controller may be configured to prevent writes from head nodes that are not authorized to write to the particular portions of the mass storage devices of the data storage sled that includes the sled controller. The sled controller may refuse to perform a write request based on being presented an invalid token or based on a token not being included with a write request.

[0048] In some embodiments, a control plane such as a local control plane or a zonal control plane of a data storage system may issue unique sequence numbers to head nodes of the data storage system to indicate which head node is a primary head node for a particular volume or volume partition. A primary head node may present a sequence number issued from a control plane to respective ones of the sled controllers of respective ones of the data storage sleds to reserve, for a particular volume or volume partition, respective portions of mass storage devices serviced by the respective ones of the respective sled controllers. In response, the sled controllers may issue a token to the primary head node to be included with future write requests directed to the respective portions.

[0049] In order to facilitate a failover operation between a primary head node and a secondary head node, a control plane may issue new credentials, e.g. a new sequence number, to a head node assuming a role of primary head node for a volume or volume partition. The newly assigned primary head node may present the credentials, e.g. new sequence number, to respective sled controllers to receive respective tokens that supersede tokens previously used to a previous head node acting as a primary head node for a particular volume or volume partition that had data stored in portions of mass storage devices service by the sled controller. Thus, during a fail over event, a previous primary head node may be fenced off from portions of mass storage devices to prevent corruption of data stored on the mass storage devices during the failover event.

[0050] FIG. 1 illustrates a data storage unit comprising head nodes and data storage sleds, according to some embodiments. Data storage unit **100**, which may be included in a data storage system, includes network switches **102** and **104**, head nodes **106** and data storage sleds **134-144** on shelves **118**. Each data storage sled **134-144** includes a sled controller **112** and mass storage devices **110**. The head nodes **106**, data storage sleds **134-144**, and network switches **102** and **104** are mounted in rack **130**. In some embodiments, networking devices, such as network switches **102** and **104**, may be mounted in a position adjacent to and external from a rack of a data storage unit, such as rack **130** of data storage unit **100**. A data storage unit may have redundant network connections to a network external to the data storage unit, such as network **128** that is connected to both network switch **102** and network switch **104**. In some embodiments, components of a data storage unit, such as network switches

102 and **104**, head nodes **106**, and data storage sleds **134-144** may be connected to redundant power sources. For example, power connections **108** indicate power connections for network switches **102** and **104**, head nodes **106**, and data storage sleds **134-144**. Note that power connections **108** are illustrated as a power symbol for simplicity of illustration, but may include various types of power connectors and power distribution systems. For example, power connectors of data storage unit components, such as head nodes and data storage sleds, may couple to dual power distribution systems within a data storage unit that receive power from dual power sources. In some embodiments, a data storage unit may include more than two redundant power distribution systems from more than two redundant power sources.

[0051] Each head node of a data storage unit, such as each of head nodes **106**, may include a local data storage and multiple network interface cards. For example, a head node may include four network ports, wherein two network ports are used for internal communications with data storage sleds of a data storage unit, such as data storage sleds **134-144**, and two of the network ports are used for external communications, for example via network **128**. In some embodiments, each head node may be assigned two publicly routable network addresses that are routable from client devices in network **128** and may also be assigned two local network addresses that are local to a data storage unit and are routable for communications between the head node and data storage sleds of the data storage unit. Thus, a data storage unit, such as data storage unit **100**, may include multiple redundant networks for communications within the data storage unit. In some embodiments, publicly routable network addresses may be used for internal communications between head nodes and data storage sleds and a head node may be assigned four publicly routable network addresses that are routable from client devices in network **128**. The data storage unit may also include redundant power distribution throughout the data storage unit. These redundancies may reduce risks of data loss or downtime due to power or network failures. Because power and network failure risks are reduced via redundant power and network systems, volumes may be placed totally or at least partially within a single data storage unit while still meeting customer requirements for reliability and data durability.

[0052] Also, one or more head nodes of a data storage unit, such as one or more of head nodes **106**, may function as a head node and additionally implement a local control plane for a data storage unit. In some embodiments, a local control plane may be implemented in a logical container separate from other control and storage elements of a head node. A local control plane of a data storage unit may select amongst any of the head nodes, such as any of head nodes **106**, of the data storage unit when selecting a head node to designate as a primary head node for a volume or volume partition and may select amongst any of the remaining head nodes of the data storage unit when selecting a head node to designate as a secondary head node for the volume or volume partition. For example a first one of head nodes **106** may be designated as a primary head node for a volume or volume partition and any of the remaining head nodes **106** may be selected as a secondary head node for the volume or volume partition. In some embodiments, a given one of the head nodes **106** may be designated as a primary head node for a given volume or volume partition and may also be designated as a secondary head node for another volume or volume partition.

[0053] Additionally, any head node may be assigned or select columns of space on mass storage devices in any of the data storage sleds of a data storage unit for storing data for a particular volume or volume partition. For example, any of head nodes **106** may reserve columns of space in mass storage devices **110** in any of data storage sleds **134-144**. However, any particular column of space of a mass storage device may only be assigned to a single volume or volume partition at a time.

[0054] Because multiple head nodes and multiple data storage sleds are available for selection, a failure of a particular head node or a failure of a mass storage device in a particular data storage sled may not significantly reduce durability of data stored in the data storage unit. This is because, upon failure of a head node, a local control plane may designate another head node of the data storage unit to function as secondary head node for a volume or volume partition. Thus, the volume is only without a secondary head node for a short period of time during which a new secondary head node is being designated and index data is being replicated from the primary head node to the secondary head node. Furthermore, when a head node of a data storage unit fails, other head nodes of the data storage unit may still be able to access data in all of the storage sleds of the data storage unit. This is because no single data storage sled is exclusively assigned to any particular head node, but instead columns of space on individual mass storage devices of the data storage sleds are assigned to particular head nodes for particular volumes or volume partitions. This arrangement greatly reduces the blast radius of a head node failure or a disk failure as compared to other storage systems in which each server has a dedicated set of storage devices.

[0055] As discussed in more detail below, in some embodiments, a head node or local control plane of a data storage unit may be configured to replicate data stored on mass storage devices that are located in a data storage sled to other mass storage devices in other data storage sleds. Thus, for example, when a data storage sled with a failed mass storage device is removed from a data storage unit for replacement or repair, data from one or more non-failed mass storage devices in a data storage sled may still be available because the data has been replicated to other data storage sleds of the data storage unit. For example, if a single mass storage device **110** in data storage sled **134** failed, data stored in the remaining mass storage devices **110** of data storage sled **134** may be replicated to mass storage devices **110** in any of data storage sleds **136-144**. Thus while data storage sled **134** is removed from data storage unit **100** for repair or replacement of the failed mass storage device **110**, data previously stored on the non-failed mass storage devices **110** of data storage sled **134** may still be available to head nodes **106**.

[0056] Also, a data storage unit, such as data storage unit **100**, may perform read and write operations independent of a zonal control plane. For example, each of head nodes **106** may be assigned one or more network addresses, such as IP addresses, that are advertised outside of data storage unit **100**. Read and write requests may be routed to individual head nodes at the assigned network addresses of the individual head nodes via networking devices of the data storage unit, such as network switches **102** and **104**, without the read and write requests being routed through a control plane external to the data storage unit, such as a control plane external to data storage unit **100**.

[0057] In some embodiments, a data storage sled, such as one of data storage sleds 134-144, may include a sled controller, such as one of sled controllers 112. A sled controller may present the mass storage devices of the data storage sled to the head nodes as storage destination targets. For example head nodes and data storage sleds may be connected over an Ethernet network. In some embodiments, head nodes, such as head nodes 106 may communicate with mass storage devices 110 and vice versa via sled controllers 112 using a Non-volatile Memory Express (NVMe) protocol, or other suitable protocols. In some embodiments, each head node may be assigned multiple private network addresses for communication with data storage sleds over redundant internal Ethernet networks internal to a data storage unit. In some embodiments, a head node at an I/O processing software layer may perform a local disk operation to write or read from a mass storage device of a data storage sled and another software layer of the head node may encapsulate or convert the I/O operation into an Ethernet communication that goes through a networking device of the data storage unit to a sled controller in one of the data storage sleds of the data storage unit. A network interface of a head node may be connected to a slot on a motherboard of the head node, such as a PCIe slot, so that the mass storage devices of the data storage sleds appears to the operating system of the head node as a local drive, such as an NVMe drive. In some embodiments, a head node may run a Linux operating system or other type of operating system. The operating system may load standard drivers, such as NVMe drivers, without having to change the drivers to communicate with the mass storage devices mounted in the data storage sleds.

[0058] In some embodiments, a local control plane may be configured to designate more than one head node as a secondary/back-up head node for a volume or a volume partition and also adjust a number of mass storage devices that make up a RAID array for longer term storage of data for the data volume or volume partition. Thus if increased durability is desired for a particular volume or volume partition, the volume data may be replicated on “N” head nodes and subsequently stored across “M” mass storage devices in data storage sleds of the data storage unit, wherein the number “N” and the number “M” may be adjusted to achieve a particular level of durability. In some embodiments, such an arrangement may allow high levels of durability to be realized without having to store data for a data volume outside of a single data storage unit. Also, in such an arrangement, input/output operations may be performed more quickly because data for a particular volume is stored within a single data storage unit.

[0059] Also, a given head node may be designated as a primary head node or a secondary head node for multiple volumes. Furthermore, a zonal control plane of a data storage system or a local control plane of a data storage unit may balance volume placement across head nodes of a data storage unit. Because volumes are distributed amongst the head nodes, variations in peak IOPS to average IOPS may be reduced because while one volume may experience peak load other volumes serviced by a particular head node may experience less than peak IOPS load. In some embodiments, a zonal or local control plane may adjust head node designations or volume assignments to balance loads if volumes on a particular head node experience significantly more IOPS than volumes serviced by other head nodes.

[0060] While, FIG. 1 illustrates mass storage devices 110 as solid state drives, any suitable storage device may be used. For example, in some embodiments, storage devices 110 may include hard disk drives. Also, FIG. 1 illustrates networking devices 102 and 104 to be networking switches. However, in some embodiments, other suitable networking devices may be used such as routers, etc.

[0061] In some embodiments, a data storage unit, such as data storage unit 100, may be part of a larger provider network system. Also, in some embodiments more than one data storage unit may be included in a block storage service of a provider network. For example, FIG. 2 illustrates such an example provider network, according to some embodiments.

[0062] FIG. 2 is a block diagram illustrating a provider network that includes multiple network-based services such as a block-based storage service that implements dynamic resource creation to connect with client resources, according to some embodiments. Provider network 200 may be set up by an entity such as a company or a public sector organization to provide one or more services (such as various types of cloud-based computing or storage) accessible via the Internet and/or other networks to clients 210. Provider network 200 may include numerous data centers hosting various resource pools, such as collections of physical and/or virtualized computer servers, storage devices, networking equipment and the like (e.g., computing device 2100 described below with regard to FIG. 21), needed to implement and distribute the infrastructure and services offered by the provider network 200. In some embodiments, provider network 200 may provide computing resources, such as virtual compute service 240, storage services, such as block-based storage service 220, and/or any other type of network-based services 260. Clients 210 may access these various services offered by provider network 200 via network 270. Likewise network-based services may themselves communicate and/or make use of one another to provide different services. For example, computing resources offered to clients 210 in units called “instances,” such as virtual or physical compute instances, may make use of particular data volumes 226, providing virtual block-based storage for the compute instances. Also, note that any of the data storage units 224a, 224b, 224n may be data storage units such as data storage unit 100 illustrated in FIG. 1.

[0063] As noted above, virtual compute service 240 may offer various compute instances, such as compute instances 254a and 254b to clients 210. A virtual compute instance may, for example, comprise one or more servers with a specified computational capacity (which may be specified by indicating the type and number of CPUs, the main memory size, and so on) and a specified software stack (e.g., a particular version of an operating system, which may in turn run on top of a hypervisor). A number of different types of computing devices may be used singly or in combination to implement the compute instances of virtual compute service 240 in different embodiments, including special purpose computer servers, storage devices, network devices and the like. In some embodiments instance clients 210 or any other user may be configured (and/or authorized) to direct network traffic to a compute instance. In various embodiments, compute instances may mount, connect, attach or map to one or more data volumes 226 provided by block-based storage service 220 in order to obtain persistent block-based storage for performing various operations.

[0064] Compute instances may operate or implement a variety of different platforms, such as application server instances, Java™ virtual machines (JVMs), special-purpose operating systems, platforms that support various interpreted or compiled programming languages such as Ruby, Perl, Python, C, C++ and the like, or high-performance computing platforms) suitable for performing client applications, without for example requiring the client **210** to access an instance.

[0065] Compute instance configurations may also include compute instances with a general or specific purpose, such as computational workloads for compute intensive applications (e.g., high-traffic web applications, ad serving, batch processing, video encoding, distributed analytics, high-energy physics, genome analysis, and computational fluid dynamics), graphics intensive workloads (e.g., game streaming, 3D application streaming, server-side graphics workloads, rendering, financial modeling, and engineering design), memory intensive workloads (e.g., high performance databases, distributed memory caches, in-memory analytics, genome assembly and analysis), and storage optimized workloads (e.g., data warehousing and cluster file systems). Size of compute instances, such as a particular number of virtual CPU cores, memory, cache, storage, as well as any other performance characteristic. Configurations of compute instances may also include their location, in a particular data center, availability zone, geographic, location, etc., and (in the case of reserved compute instances) reservation term length.

[0066] As illustrated in FIG. 2, a virtualization host, such as virtualization hosts **242a** and **242n**, may implement and/or manage multiple compute instances **252a**, **252b**, **254a**, and **254b** respectively, in some embodiments, and may be one or more computing devices, such as computing device **2100** described below with regard to FIG. 21. Virtualization hosts **242** may also provide multi-tenant hosting of compute instances. For example, in some embodiments, one virtualization host may host a compute instance for one entity (e.g., a particular client or account of virtual computing service **210**), while another compute instance hosted at the same virtualization host may be hosted for another entity (e.g., a different account). A virtualization host may include a virtualization management module, such as virtualization management modules **244a** and **244b** capable of instantiating and managing a number of different client-accessible virtual machines or compute instances. The virtualization management module may include, for example, a hypervisor and an administrative instance of an operating system, which may be termed a “domain-zero” or “dom0” operating system in some implementations. The dom0 operating system may not be accessible by clients on whose behalf the compute instances run, but may instead be responsible for various administrative or control-plane operations of the network provider, including handling the network traffic directed to or from the compute instances.

[0067] Virtual computing service **240** may implement control plane **250** to perform various management operations. For instance, control plane **250** may implement resource management to place compute instances, and manage the access to, capacity of, mappings to, and other control or direction of compute instances offered by provider network. Control plane **250** may also offer and/or implement a flexible set of resource reservation, control and access interfaces for clients **210** via an interface (e.g., API). For

example, control plane **250** may provide credentials or permissions to clients **210** such that compute instance control operations/interactions between clients and in-use computing resources may be performed.

[0068] In various embodiments, control plane **250** may track the consumption of various computing instances consumed for different virtual computer resources, clients, user accounts, and/or specific instances. In at least some embodiments, control plane **250** may implement various administrative actions to stop, heal, manage, or otherwise respond to various different scenarios in the fleet of virtualization hosts **242** and instances **252**, **254**. Control plane **250** may also provide access to various metric data for client(s) **210** as well as manage client configured alarms.

[0069] In various embodiments, provider network **200** may also implement block-based storage service **220** for performing storage operations. Block-based storage service **220** is a storage system, composed of one or more computing devices implementing a zonal control plane **230** and a pool of multiple data storage units **224a**, **224b** through **224n** (e.g., data storage units such as data storage unit **100** illustrated in FIG. 1), which provide block level storage for storing one or more sets of data volume(s) **226a**, **226b** through **226n**. Data volumes **226** may be attached, mounted, mapped, or otherwise connected to particular clients (e.g., a virtual compute instance of virtual compute service **240**), providing virtual block-based storage (e.g., hard disk storage or other persistent storage) as a contiguous set of logical blocks. In some embodiments, a data volume **226** may be divided up into multiple data chunks or partitions (including one or more data blocks) for performing other block storage operations, such as snapshot operations or replication operations. A volume snapshot of a data volume **226** may be a fixed point-in-time representation of the state of the data volume **226**. In some embodiments, volume snapshots may be stored remotely from a data storage unit **224** maintaining a data volume, such as in another storage service **260**. Snapshot operations may be performed to send, copy, and/or otherwise preserve the snapshot of a given data volume in another storage location, such as a remote snapshot data store in other storage service **260**. In some embodiments, a block-based storage service, such as block-based storage service **220**, may store snapshots of data volumes stored in the block-based storage service.

[0070] Block-based storage service **220** may implement zonal control plane **230** to assist in the operation of block-based storage service **220**. In various embodiments, zonal control plane **230** assists in creating volumes on data storage units **224a**, **224b**, through **224n** and moving volumes between data storage units **224a**, **224b**, through **224n**. In some embodiments, access to data volumes **226** may be provided over an internal network within provider network **200** or externally via network **270**, in response to block data transaction instructions.

[0071] Zonal control plane **230** may provide a variety of services related to providing block level storage functionality, including the management of user accounts (e.g., creation, deletion, billing, collection of payment, etc.). Zonal control plane **230** may implement capacity management, which may generate and manage a capacity model for storage service **220**, and may direct the creation of new volumes on particular data storage units based on the capacity of storage service **220**. Zonal control plane **230** may

further provide services related to the creation and deletion of data volumes 226 in response to configuration requests.

[0072] Clients 210 may encompass any type of client configured to submit requests to network provider 200. For example, a given client 210 may include a suitable version of a web browser, or may include a plug-in module or other type of code module configured to execute as an extension to or within an execution environment provided by a web browser. Alternatively, a client 210 may encompass an application such as a database application (or user interface thereof), a media application, an office application or any other application that may make use of compute instances, a data volume 226, or other network-based service in provider network 200 to perform various operations. In some embodiments, such an application may include sufficient protocol support (e.g., for a suitable version of Hypertext Transfer Protocol (HTTP)) for generating and processing network-based services requests without necessarily implementing full browser support for all types of network-based data. In some embodiments, clients 210 may be configured to generate network-based services requests according to a Representational State Transfer (REST)-style network-based services architecture, a document- or message-based network-based services architecture, or another suitable network-based services architecture. In some embodiments, a client 210 (e.g., a computational client) may be configured to provide access to a compute instance or data volume 226 in a manner that is transparent to applications implemented on the client 210 utilizing computational resources provided by the compute instance or block storage provided by the data volume 226.

[0073] Clients 210 may convey network-based services requests to provider network 200 via external network 270. In various embodiments, external network 270 may encompass any suitable combination of networking hardware and protocols necessary to establish network-based communications between clients 210 and provider network 200. For example, a network 270 may generally encompass the various telecommunications networks and service providers that collectively implement the Internet. A network 270 may also include private networks such as local area networks (LANs) or wide area networks (WANs) as well as public or private wireless networks. For example, both a given client 210 and provider network 200 may be respectively provisioned within enterprises having their own internal networks. In such an embodiment, a network 270 may include the hardware (e.g., modems, routers, switches, load balancers, proxy servers, etc.) and software (e.g., protocol stacks, accounting software, firewall/security software, etc.) necessary to establish a networking link between given client 210 and the Internet as well as between the Internet and provider network 200. It is noted that in some embodiments, clients 210 may communicate with provider network 200 using a private network rather than the public Internet.

Data Replication

[0074] FIG. 3 is a block diagram illustrating head nodes and data storage sleds of a data storage unit storing block storage data in response to a write request, according to some embodiments. Head nodes 306 illustrated in FIG. 3 may be the same as head nodes 106 illustrated in FIG. 1. Also, data storage sleds 326 may be the same as data storage sleds 134-144 illustrated in FIG. 1.

[0075] As discussed above, a data storage system that includes a data storage unit, may store volume data in a data storage of a first head node designated as a primary head node for a volume or volume partition and may also replicate the volume data to a second head node designated as a secondary head node for the volume or volume partition. For example, at time 1, a write request 302 is routed to head node 306 that is designated as a primary head node for a volume or volume partition. At time 2 subsequent to the write request being received at head node 306, data included with the write request is stored in storage 314 of primary head node 306 and primary head node 306 causes the data included with the write request to be replicated to storage 316 of secondary head node 308. Replication of the data to secondary head node 306 is performed concurrently or nearly concurrently with storing the data in storage 314 of primary head node 306. Also, as shown in FIG. 3 at time 2, replication of the data to the secondary head node may include the secondary head node sending an acknowledgment back to the primary head node indicating that the data has been replicated to the secondary head node. Subsequently at time 3, which is also nearly concurrent with the data being stored in the storage of the primary head node and the data being replicated to the secondary head node, the primary head node, head node 306, may issue an acknowledgement 320 to the client device that requested write 302 has been committed in data storage system 300.

[0076] In some embodiments, a write request, such as write request 302, may be concurrently received at a primary head node and a secondary head node. In such embodiments, the primary head node may verify that the secondary head node has committed the write before acknowledging at time 3 that the write has been committed in the data storage system.

[0077] At a later point in time 4, e.g. asynchronous to times 1-3, the primary head node, e.g. head node 306, may cause data stored in storage 314, that includes the data included with the write request and that may include additional data stored before or after the write request, to be flushed to mass storage devices 322 of the data storage sleds 326 of the data storage unit. For example, at time 4 data is flushed to mass storage devices 322 of data storage sleds 326. In some embodiments, data is divided into portions and stored across multiple mass storage devices, each in a different sled and/or on a different shelf of a data storage unit. In some embodiments, data is also erasure encoded when stored in mass storage devices of data storage sleds. For example, data flushed from storage 314 of head node 306 may be divided into six portions where each portion is stored in a different mass storage device of a different data storage sled on a different shelf of a data storage unit 350 of data storage system 300 and is also erasure encoded across the different mass storage devices. For example data portions are stored in sled A of shelf 1, sled B of shelf 2, sled A of shelf 3, sled C of shelf 4, sled B of shelf 5, and sled C of shelf 6.

[0078] Also, as can be seen in FIG. 3, a data storage unit, such as data storage unit 350, may include “M” number of shelves and “N” number of head nodes. The portions of data may be stored on portions of mass storage devices 322 in the respective data storage sleds 326. In order to distinguish between a portion of data and a portion of space on a mass storage device, a portion of space on a mass storage device may be referred to herein as a “column” of a mass storage

device. Furthermore, a set of columns of mass storage devices that store different portions of data of a volume such as the columns shown in sled A of shelf 1, sled B of shelf 2, sled A of shelf 3, sled C of shelf 4, sled B of shelf 5, and sled C of shelf 6 may collectively make up what is referred to herein as an “extent.” For example, in an erasure encoded RAID six array, an extent may include six columns that collectively make up the RAID array. Four of the columns may store striped data and two of the columns may store parity data. In some embodiments, other replication algorithms other than erasure encoding may be used such as quorum algorithms, etc.

[0079] In some embodiments, each column of an extent may be in a different fault domain of a data storage unit. For example, for the extent being stored in FIG. 3 each column is located in a different data storage sled that is mounted on a different shelf of the data storage unit 350. Thus a failure of a sled controller, such as one of sled controllers 324, may only affect a single column. Also if a power supply of a data storage sled fails it may only affect a single data storage sled or if a part of a power distribution system fails it may affect a single shelf. However, because each column of an extent may be located in a different shelf, a shelf level power event may only affect a single column of the extent.

[0080] In some embodiments, a head node of a data storage unit, such as one of head nodes 304, 306, 308, or 310, may implement a local control plane. The local control plane may further implement an extent allocation service that allocates extents to head nodes designated as a primary head node for a volume or volume partition. In some embodiments, an extent allocation service may allocate a set of extents to a particular volume referred to herein as a “sandbox.” The primary head node for the particular volume may then select extents to store data on during a data flush from the primary head node to data storage sleds of the data storage unit by selecting an extent from the sandbox allocated for the particular volume.

[0081] In some embodiments, if insufficient space is available in the particular volume’s sandbox or if a particular placement would cause a data durability of data to be saved to fall below a minimum required durability for the particular volume, a primary head node for the particular volume may select columns outside of the particular volume’s sandbox to write data for the particular volume. For example, a sandbox may include multiple columns that make up multiple extents in different ones of the data storage sleds 326 on different ones of the shelves of a data storage unit 350. A primary head node may be able to flush data to columns within a particular volume’s sandbox without having to request extent allocation from a local control plane that implements an extent allocation service. This may further add durability and reliability to a data storage unit because a primary head node for the particular volume may continue to flush data even if communication is lost with a local control plane within the data storage unit. However, if space is not available or a placement would cause durability for a particular volume or volume partition to fall below a minimum threshold, a primary head node may flush data to columns outside of the particular volume’s sandbox. In some embodiments, a primary head for a particular volume may flush data to columns outside the primary head node’s sandbox without requesting an allocation from a local control plane that implements an extent allocation service. For example, a primary head node may store addresses for each

sled controller in a data storage unit and may flush data to any sled controller in the data storage unit that is associated with mass storage devices with available columns.

[0082] As will be discussed in more detail in regard to FIG. 15, a sled controller of a data storage sled, such as sled controller 324, may implement a fencing protocol that prevents a primary head node from writing to columns for which another primary head node has assumed control after the primary head node has been superseded by another head node assuming the role of primary head node for a particular volume or volume partition. It should be pointed out that a secondary head node or other back-up head nodes may not flush data to data storage sleds and flushing may be limited to only being performed by a primary head node.

[0083] Because for a particular volume, the volume’s data may be stored in a storage of a primary head node and replicated to a secondary head node and may later be moved to being stored across an extent of mass storage devices in different data storage sleds of a data storage unit, an index with pointers to where the data is stored may be used for subsequent read requests and write requests to locate the data. Also in some embodiments, storages of a head node may be log-structured such that incoming write request are written to the head of the log of the head node’s log-structured storage. An index entry may be added indicating where the written data is stored in the head node’s log and subsequently the index may be updated when the written data is flushed from the log of the primary head node to an extent comprising columns of mass storage devices of the data storage system.

[0084] FIGS. 4A-4B are block diagrams illustrating a log-structured storage and an index of a head node storage, according to some embodiments. Head node 402 includes storage 404 that includes log 408 and index 406. Volume data may be stored in log 408 prior to being flushed to mass storage devices of a data storage unit. Index information 410 may include an entry for the volume data and a corresponding pointer to where the volume data is stored. For example, index information 410 indicates that data for volume 1, offset A, length B is stored in log storage 408 at log segment C and offset D. In some embodiments, a log of a head node such as log 408 of storage 404 of head node 402 may store data for more than one volume. For example, index information 410 also includes an entry for volume 2 offset E, length F and a corresponding pointer indicating the data for this volume entry is stored in log 408 at log segment G, offset H.

[0085] While FIGS. 4A-B illustrate log storage 408 and index 406 as separate from each other, in some embodiments, an index, such as index 406, may lay on top of a log or side-by-side with a log, such as log storage 408.

[0086] When data for a volume is moved from a storage of a head node to being stored in an extent across multiple mass storage devices of a data storage unit, the data for the volume may be removed from a log of a head node storage and an index of the head node storage may be updated to indicate the new location at which the data for the volume is stored. For example, in FIG. 4B, index information 412 indicates that data for volume 1, offset A, length B is now stored at extent A and data for volume 2, offset E, length F is now stored at extent B. Note that the labels “extent A” and “extent B” are used for ease of illustration. In some embodiments, an index may include addresses of data storage sleds where the data for the volume is located, such as local IP

addresses of the data storage sleds, and addresses of the columns of the mass storage devices within the data storage sleds. In some embodiments, an index may include another label such as “extent A” where each head node stores information for locating “extent A” or may consult an extent allocation service for locating “extent A.” In some embodiments, an index may include addresses of data storage sleds where the data for the volume is located and sled controllers of the data storage sleds may be able to determine the appropriate columns based on volume IDs stored in respective columns allocated to the volume.

[0087] When a read request is received by a head node designated as a primary head node for a volume, the head node may consult an index of a storage of the head node, such as index 406 of storage 404, to determine what is the latest version of the volume’s data and where the latest version of the volume’s data is stored. For example a primary head node, such as head node 402, may consult the primary head node’s index, such as index 406, to determine if the latest version of the volume’s data is stored in the head node’s log, such as log 408, or is stored in an extent comprising mass storage devices of the data storage unit.

[0088] FIG. 5 illustrates a partial view of a data storage unit that stores portions of a volume partition in multiple mass storage devices in multiple data storage sleds on multiple shelves of the data storage unit, according to some embodiments. FIG. 5 illustrates an example storage pattern for extent A from index 406 in FIG. 4B. Extent A from index 406 illustrated in FIG. 4B is shown as extent A 502 in FIG. 5. Also, an example storage pattern for extent B from index 406 illustrated in FIG. 4B is shown in FIG. 5 as extent B 504. Note that a data storage sled may include multiple columns of multiple extents. Also, in some embodiments a single mass storage device may include multiple columns of multiple extents.

[0089] FIGS. 6A-B illustrate columns of mass storage devices storing different portions of a volume partition, according to some embodiments. FIG. 6A illustrates an embodiment in which data flushed to extent A, which may be the same extent A as described in FIGS. 4 and 5, is erasure encoded across 4+2 columns. The striped data 602 may include the original data flushed from log 408 divided into multiple portions and the parity data 604 may include encoded data that allows the flushed data to be recreated in case of failure of one or more of the mass storage devices or sleds that include one of the columns. FIG. 6B illustrates a similar embodiment where extent B is erasure encoded across four striped data columns 606 and two parity columns 608. Note that in FIG. 6B the data is stored in a different location in the column than is shown in FIG. 6A. This is intended to illustrate that the columns shown in FIG. 6B may already store data previously written to the columns of extent B, whereas the data being written to extent A may be the first set of data written to extent A. Also, it is worth noting that for a particular volume, multiple extents may be assigned to store data of the volume. In some embodiments, an extent may represent a fixed amount of storage space across a set number of columns of mass storage devices. When an extent is filled for a particular volume, another extent may be allocated to the volume by a head node or an extent allocation service. FIGS. 6A and 6B illustrate an example RAID level and erasure encoding technique. However, in some embodiments various other RAID levels may be used and various data coding techniques may be used to

increase durability of stored data. It also worth noting that erasure encoding data may reduce a number of columns needed to achieve a particular level of durability. For example, data stored that is not erasure encoded may require the data to be stored redundantly across 8 columns to achieve a given level of durability, whereas a similar level of durability may be achieved by erasure encoding the data across fewer columns. Thus erasure encoding data may significantly reduce an amount of storage resources that are needed to store data to a particular level of durability. For example, data erasure encoded according to a 4+2 erasure coding scheme may be recreated from any four of the six columns, wherein the six columns include four columns of striped data segments and two columns of parity data segments.

[0090] In some embodiments, a data storage system may implement one or more communication protocols between head nodes and data storage sleds of the data storage system that allow for rapid communications between the head nodes and the data storage sleds. Thus, high levels of performance may be provided to clients of a data storage system despite volume data being erasure encoded across multiple columns of mass storage devices in different data storage sleds. For example, a data storage system may implement a protocol for reliable out-of-order transmission of packets as described in U.S. patent application Ser. No. 14/983,436 filed on Dec. 29, 2015, which is herein incorporated by reference. Also, for example, a data storage system may implement a protocol for establishing communication between a user application and a target application wherein the network does not require an explicit connection between the user application and the target application as described in U.S. patent application Ser. No. 14/983,431 filed on Dec. 29, 2015, which is herein incorporated by reference. In some embodiments, implementation of such protocols may permit data erasure encoded across multiple mass storage devices in multiple different data storage sleds to be read by a head node in a timely manner such that, from a perspective of a client device of the data storage system, performance is comparable to a system that does not erasure encode volume data across multiple mass storage devices or such that performance exceeds a performance of a system that does not erasure encode volume data across multiple mass storage devices.

[0091] FIG. 7 is a high-level flowchart illustrating operations performed by a head node in response to a request to store data in a data storage unit, according to some embodiments.

[0092] At 702, upon receiving a write request from a client device, wherein the write request is directed to a particular volume for which the head node is functioning as a primary head node, the head node writes data included with the write request to the log of the head node and updates the index of the head node to include an entry for the volume data and a pointer indicating where the volume data is stored.

[0093] At 704, the primary head node causes the data included with the write request to be replicated to the secondary head node. The secondary head node then stores the data in a log of the secondary head node and updates an index of a storage of the secondary head node to include an entry for the volume data and a pointer indicating where the volume data is stored. The secondary head node may then send an acknowledgement to the primary head node indicating that the data has been replicated in the secondary head

node's storage. In some embodiments, the primary head node then issues an acknowledgement to the client device indicating that the requested write has been persisted in the data storage system. In some embodiments, replication between head nodes could be primary and secondary e.g. master/slave replication. In some embodiments, other replication techniques such as a Paxos protocol, other consensus protocol, etc. may be used to replicate data between head nodes.

[0094] At **706**, the primary head node determines if the log data of the primary head node exceeds a threshold that would trigger the log data or a segment of the primary head node's log data to be flushed to extents that include columns of mass storage devices of data storage sleds of a data storage unit that includes the head node. In some embodiments, a threshold to trigger data to be flushed may include: an amount of data stored in the log or in a segment of the log, an amount of time that has elapsed since the data was last accessed or altered, a frequency at which the data is accessed or altered, or other suitable thresholds. In some embodiments, data flushed from a log of a head node may only include a portion of the data written to the log of the head node or a segment of the log of the head node. For example, older data stored in a log of a head node may be flushed while more recently written data may remain in the log of the head node. In some embodiments, a frequency of flush operations from a log of a head node may be throttled based on a variety of factors, such as a fill rate of the log of the head node or based on an amount of write requests being received by the head node or being received for a particular volume serviced by the head node.

[0095] In response to determining the threshold has not been met, the primary head node continues to write data to the log and reverts to **702**.

[0096] At **708**, in response to determining that the threshold has been met or exceeded, the primary head node causes data stored in the log of the primary head node or a segment of the log of the primary head node to be flushed to columns of mass storage devices in different ones of a plurality of data storage sleds of the data storage unit.

[0097] At **710**, the primary head node updates the log of the primary head node to include a pointer for the volume data indicating that the flushed volume data is now stored in particular columns of mass storage devices or an extent that includes multiple columns of mass storage devices.

[0098] At **712**, the primary head node causes the secondary head node to update an index of the secondary head node to indicate the new location of the volume data. The secondary head node also releases the log space in the secondary head node that previously stored the replicated volume data.

[0099] At **714**, the head node acting as primary head node also releases space in the primary head node's log. In some embodiments, a garbage collection mechanism may cause log space to be released based on inspecting an index of a storage of a head node. In some embodiments, releasing log storage space may be performed concurrently with flushing log data or may be performed at some time subsequent to flushing log data.

[0100] FIG. **8A** is a high-level flowchart illustrating operations performed by a head node in response to a failed mass storage device in a data storage sled of a data storage unit, according to some embodiments.

[0101] At **802**, a head node or a sled controller detects a failed mass storage device in a particular data storage sled. For example, a data storage sled may include multiple mass storage devices, such as solid state storage drives, and one of the mass storage devices may fail. In some embodiments, a data storage sled may include disk drives and one of the disk drives may fail. In some embodiments, a data storage sled may include other types of mass storage devices.

[0102] At **804**, a head node acting as a primary head node for a volume with extents that include one or more columns on the failed mass storage device or a local control plane for the data storage unit causes the extents that include columns on the failed mass storage device to be replicated to other extents that include columns on other mass storage devices in other sleds of the data storage unit. For example, in a 4+2 erasure coding scheme data from any one lost mass storage drive can be recreated based on data stored on the other mass storage devices that make up an extent. Thus, data previously stored on the failed mass storage device can be recreated and replicated to data storage sleds that do not include a failed mass storage device.

[0103] At **806**, indexes of a primary head node and a secondary head node that are designated for each volume that included an extent in the failed mass storage device are updated to indicate the new locations of the data for the volumes.

[0104] In some embodiments, a data storage system may continue to operate a data storage sled that includes a failed mass storage device, such as the failed mass storage device at **808**. In some embodiments, step **806** may be omitted and all extents stored on mass storage devices in the data storage sled that includes the failed mass storage device may be replicated to other data storage sleds. Because the extents that include columns on the failed mass storage device have been replicated to data storage sleds that do not include failed mass storage devices, the durability of the data previously stored on the failed mass storage device has been recovered to the original level of durability. For example in a RAID configuration of six segments, the number of segments is returned to six by replicating the data from the failed mass storage device to other mass storage devices in the data storage unit.

[0105] FIG. **8B** is a high-level flowchart illustrating operations performed by a head node in response to a failed mass storage device in a data storage sled of a data storage unit, according to some embodiments.

[0106] In some embodiments, a data storage system may tolerate one or more failed mass storage devices in a particular sled before the mass storage devices are replaced. For example, at **852** one or more additional failed mass storage devices are detected in a data storage sled. In some embodiments the additional failed mass storage devices may be in the same data storage sled as the failed mass storage device described in FIG. **8A** or may be in a different data storage sled of the data storage unit.

[0107] At **854**, data from other non-failed mass storage devices each in a data storage sled that includes a failed mass storage device is copied to other mass storage devices in other data storage sleds of the data storage unit. In some embodiments, only data from non-failed mass storage devices that are included in a data storage sled that is to be repaired may be copied. In some embodiments, copying the data from the non-failed mass storage devices may include recreating the data from a set of columns stored on remain-

ing non-failed mass storage devices and then erasure encoding the data across another set of columns of mass storage devices of a replacement extent. For example, in a 4+2 erasure encoding scheme, data of an extent may be recreated from any four of the six columns of the extent. After being recreated, the data may be erasure encoded across another set of 4+2 columns of a replacement extent.

[0108] At 856, indexes of a primary head node and a secondary head node that are designated for each volume that included an extent in the affected mass storage devices are updated to indicate the new locations of the data for the volumes that has been copied to other mass storage devices in the data storage unit.

[0109] At 858, the data storage sled(s) that includes the failed mass storage device is at least partially removed from the data storage unit and the failed mass storage device is replaced. Because data previously stored on the non-failed mass storage devices of the data storage sled being removed has been copied to other mass storage devices of the data storage unit, the data remains available even while the data storage sled is at least partially removed from the data storage unit.

[0110] At 860, the data storage sled with the replaced mass storage device is re-installed in the data storage unit. At 862 mass storage devices of the replaced data storage sled are made available for allocation of columns on the mass storage devices of the data storage sled. In some embodiments, data storage space of the non-failed mass storage devices of the data storage sled may be released and made available to store data for newly allocated extents. In some embodiments, the non-failed mass storage devices may still store volume data that has been copied to other mass storage devices in the data storage unit. In some embodiments, the indexes of the respective head nodes may be updated to indicate volume data that is still stored on the non-failed mass storage devices.

Multi-Tier Control Plane

[0111] In some embodiments, a data storage system may include multiple data storage units. Management of the data storage system may be performed by a multi-tiered control plane. For example, in some embodiments a zonal control plane may determine which data storage units new volumes are to be allocated to and may perform migration of volumes between data storage units to balance loads. Also, in some embodiments, a local control plane of a data storage unit may determine which head nodes of the data storage unit are to be assigned to a particular volume or volume partition as a primary head node and a secondary head node. Also, a local control plane may manage allocation of extents within a data storage unit via a “sandbox” technique and may perform fail over operations in response to a failure of a head node, a mass storage device, or a data storage sled. In some embodiments, a data storage unit may operate autonomously from a zonal control plane subsequent to a volume being assigned to the data storage unit. Because data storage units may operate autonomous from a zonal control plane, a failure of a zonal control plane may not impact a data storage unit’s ability to respond to read and write requests or perform fail-over operations in response to a failure of a head node or a mass storage device. Also, because a local control plane of a data storage unit only affects a single data storage unit, a failure of a local control plane may have a blast radius that is limited to a single data storage unit.

Furthermore, a data storage unit may implement a local control plane on one or more head nodes of a data storage unit and implement a lease protocol to allow for fail over of the local control plane from one head node to another head node in response to a failure of a head node implementing the local control plane. In some embodiments, a local control plane may utilize a distributed value store that is distributed across the plurality of head nodes of the data storage unit. Thus, when a particular head node implementing a local control plane fails, another head node taking over implementation of the local control plane may utilize the distributed value store without values in the value store being lost due to the failure of the head node previously implementing the local control plane.

[0112] FIG. 9A is a block diagram illustrating a process for creating a volume involving a zonal control plane, a local control plane, and head nodes of a data storage system, according to some embodiments. Data storage system 900 includes one or more computing devices that implement zonal control plane 904 and also includes data storage units 906, 928, and 930. Data storage units 906, 928, and 930 may be the same as any of the data storage units described in FIGS. 1-8. Data storage unit 906 includes head nodes 908, 912, and 914 and data storage sleds 916. A local control plane 910 is implemented on head node 908. Data storage unit 928 includes head nodes 918, 920, and 924 and data storage sleds 926. A local control plane 922 for data storage unit 928 is implemented on head node 920. Data storage unit 930 includes head nodes 932, 934, and 936 and sleds 940. A local control plane 938 for data storage unit 930 is implemented on head node 936. As can be seen a local control plane for a data storage unit can be implemented on any one or more head nodes of a plurality of head nodes of a data storage unit. In some embodiments, a local control plane may be logically separated from a data plane of a head node, for example a local control plane may be located in a separate container. In some embodiments, each head node of a data storage unit may include logically isolated program instructions for implementing a local control plane and a portion of a distributed value store distributed across logically isolated portions of respective ones of the head nodes. In such embodiments, a given head node holding a lease for implementing the local control plane may implement the local control plane using the program instructions stored in the given head node. Upon failure of the given head node, another given head node may assume the lease for implementing the local control plane and may implement the local control plane using the program instructions for implementing the local control plane stored in the other given head node. For example, the given head node and the other given head node may both store program instructions for implementing the local control plane and a single one of the given head node or the other given head node may implement the local control plane at a given time.

[0113] Client device(s) 902 may be part of a separate network that is separate from data storage system 900, such as a customer network, or may be client devices within a provider network that utilizes data storage system 900. Client device(s) 902 send volume request A 942 and volume request B 944 to zonal control plane 904 to request volumes of data storage system 900 be allocated to the client devices. In response, zonal control plane 904 issues a volume creation instruction A 946 to data storage unit 906 and a volume creation instruction B 948 to data storage unit 928. In some

embodiments, volume creation instructions from a zonal control plane may be processed by a local control plane of a data storage unit. For example, local control plane **910** of data storage unit **906** processes volume creation instruction A **946** and local control plane **922** of data storage unit **928** processes volume creation instruction B **948**. In some embodiments, a zonal control plane may receive accumulated performance and usage metrics from data storage units and assign volumes, based on the accumulated performance and usage metrics. For example, a zonal control plane may attempt to balance loads between data storage units by selecting to assign new volumes to data storage units that have accumulated performance and usage metrics that indicate less load than other data storage units.

[0114] In order to process a volume creation instruction, a local control plane may assign a head node of a data storage unit to function as a primary head node for a volume and may assign another head node of a data storage unit to function as a secondary head node for the volume. For example, local control plane **910** assigns head node **912** as a primary head node for the newly created volume via assignment **950** and assigns head node **914** as secondary head node for the volume via assignment **952**. Also, local control plane **922** of data storage unit **928** assigns head node **918** as a primary head node for a newly created volume via assignment **956** and assigns head node **924** as a secondary head node for the volume via assignment **958**. As can be seen, any one of the head nodes of a data storage unit may be selected to function as a primary or secondary head node for a given volume or volume partition. Also, a local control plane of a data storage unit may collect performance information from head nodes and select primary and secondary head nodes for a given volume based on a current loading of head nodes in a data storage unit. In some embodiments, a local control plane may attempt to balance loads between head nodes when assigning primary and secondary head nodes for a given volume.

[0115] FIG. 9B is a block diagram illustrating head nodes of a data storage unit servicing read and write requests independent of a zonal control plane of a data storage system, according to some embodiments. Data storage system **900** illustrated in FIG. 9B is the same data storage system **900** as illustrated in FIG. 9A and shows read and write requests being sent to head nodes after primary and secondary head nodes for the newly created volumes have been assigned. As can be seen, read and write requests may be serviced by data storage units of a data storage system independent of a zonal control plane. Also, in some embodiments, read and write requests may be serviced independent of a local control plane. For example read request **968** and write request **960** are directed directly to head node **912** that functions as a primary head node for the newly created volume. Also, primary head node sends read data **972** and write acknowledgement **964** to client device(s) **902** without passing the read data **972** or the write acknowledgement **964** through zonal control plane **904** or local control plane **910**. In some embodiments, each head node may be assigned at least one public network address, such as a public IP address. When a head node is assigned to function as a primary head node for a volume of a client, the primary head node's public IP address may be communicated to the client device and the client device's address may be communicated to the head node functioning as primary head node. Thus, subsequent communications between the head node and the

client device may be directly routed between the exchanged address information. In some embodiments, read and write requests may be directed to both a primary and secondary head node. For example, if a client does not receive a response from a primary head node, the client may direct a read or write request to a secondary head node. In some embodiments, this may trigger a secondary head node to attempt to assume the role of primary head node.

[0116] FIG. 10A is a block diagram of a head node, according to some embodiments. Head node **1000** may be any of the head nodes described in FIG. 1-9 or 11-19. Head node **1000** includes a data control plane **1002**, storage **1010**, local control plane **1004**, and monitoring module **1016**. A data control plane of a head node, such as data control plane **1002**, may service read and write requests directed to the head node. For example, a data control plane may store one or more public IP addresses of the head node and provide the public IP addresses of the head node to client devices to allow the client devices to communicate with the head node. A storage of a head node, such as storage **1010**, may include a log, such as log **1012**, and an index, such as index **1014**. Log **1012** and index **1014** may be similar to log **408** and index **406** as described in regard to FIGS. 4A and 4B and may store pointers for volume data indicating where the volume data is stored. In some embodiments, a data control plane, such as data control plane **1002**, may consult an index, such as index **1014**, in order to service read and write requests directed at a particular volume for which the head node is functioning as a primary head node. In some embodiments, an index, such as index **1014** may indicate whether a portion of volume data for a volume is stored in a log of the head node, such as log **1012**, or is stored in an extent across multiple data storage sleds, such as mass storage devices **1022** of data storage sled **1018** illustrated in FIG. 10B that also includes sled controller **1020**. In addition, a head node may include program instructions for implementing a local control plane that are logically isolated from the data control plane of the head node.

[0117] In some embodiments, a local control plane includes an extent allocation service, such as extent allocation service **1006**, and a distributed value store, such as value store **1008**. An extent allocation service may provide "sandbox" recommendations to head nodes of a data storage unit that include sets of columns from which the head nodes may select new extents. A value store may store extent allocation information and may also store head node assignment information. In some embodiments, a local control plane may provide sequence numbers to newly assigned primary head nodes. In some embodiments, a distributed value store, such as value store **1008**, may be implemented over all or a portion of the head nodes of a data storage unit. This may provide fault tolerance such that if any one or more of the head nodes fail, the remaining head nodes may include data from the distributed data store, such that data from the distributed data store is not lost due to the failure of the one or more head nodes.

[0118] In some embodiments, a head node includes a monitoring module, such as monitoring module **1016**. Monitoring module may collect performance and/or usage metrics for the head node. A head node, such as head node **1000** may provide performance and/or usage metrics to a local control plane, such as local control plane **1004**, or may provide performance and/or usage metrics to a zonal control plane.

[0119] FIG. 11 is a high-level flowchart illustrating a process of creating a volume in a data storage system, according to some embodiments.

[0120] At 1102, a local control plane of a data storage unit of a data storage system receives a volume assignment from a zonal control plane of the data storage system.

[0121] At 1104, the local control plane assigns a first head node of the data storage unit to function as a primary head node for the newly created or newly assigned volume. At 1106, the local control plane assigns a second head node of the data storage unit to function as a secondary head node for the newly created or newly assigned volume. Note that in some embodiments, a zonal control plane may move volume between data storage units of a data storage system. Thus the newly assigned volume may be an existing volume being moved from another data storage unit of the data storage system. Also, a local control plane of a data storage unit may select head nodes to function as primary and secondary head nodes from any of the head nodes of the data storage unit. However, a head node functioning as a primary head node may not function as a secondary head node for the same volume. But, a given head node may function as a primary head node for more than one volume and may also function as a secondary head node for one or more other volumes.

[0122] At 1108 the primary head node for the volume services read and write requests directed at the volume. In some embodiments, a head node functioning as a primary head node may service read and write requests independent of a zonal control plane and/or independent of a local control plane of a data storage unit.

[0123] FIG. 12A is a high-level flowchart illustrating a local control plane of a data storage unit providing storage recommendations to a head node of the data storage unit for locations to store data in data storage sleds of the data storage unit for a given volume, according to some embodiments.

[0124] At 1202, a local control plane of a data storage unit allocates a “sandbox” to a particular volume serviced by a primary head node functioning as primary head node for the particular volume. The sandbox may include a set of columns of mass storage devices from which the head node is recommended to select extents for the particular volume. In some embodiments, the sandbox may include extents that already include corresponding columns in multiple mass storage devices and the head node may be recommended to select extents for the particular volume from the extents included in the sandbox recommendation.

[0125] At 1204, the local control plane collects performance metrics from data storage sleds and/or head nodes in the data storage unit.

[0126] At 1206, the local control plane issues “sandbox” updates to the primary head node functioning as a primary head node for the particular volume. The sandbox updates may be based on the collected performance metrics collected at 1204. A local control plane may allocate sandbox recommendations and update sandbox recommendations to avoid heat collisions wherein multiple head nodes are attempting to access the same data storage sleds at the same time. In some embodiments, a sandbox recommendation may be a loose constraint and a head node functioning as a primary head node may select columns or extents that are not included in a sandbox recommendation. It should also be noted that sandbox recommendation and performance and/or usage metrics collection may be performed outside of the

I/O path. Thus, if there is a failure or corruption of the local control plane, reads and writes may continue to be processed by non-affected head nodes of a data storage unit. Also, a sandbox allocated to a particular volume may remain with the particular volume during a failover of head nodes. For example, if a primary head node for a particular volume fails, the sandbox allocated for the particular volume may move with the particular volume that will now be serviced by a former secondary head node. Subsequent to a head node failover, sandbox updates, such as the sandbox updates described at 1206, may be issued from the local control plane to the new primary head node for the volume.

[0127] FIG. 12B is a high-level flowchart illustrating a head node of a data storage unit storing data in data storage sleds of the data storage unit, according to some embodiments.

[0128] At 1252, a primary head node determines a segment of data to be flushed to mass storage devices in data storage sleds of a data storage unit. For example, exceeding one or more thresholds, such as an amount of data stored in a log, an age of data stored in a log, or an infrequency at which the data is accessed in a log, may trigger a primary head node to flush data to data storage sleds.

[0129] At 1254, a primary head node may determine if there is available space in a sandbox allocated to a volume serviced by the primary head node. At 1256, in response to determining there is sufficient space in the sandbox, the primary head node flushes the data to extents that include columns in the allocated sandbox allocated for the volume. At 1258, in response to determining there is insufficient space in the sandbox or in response to determining a placement in the sandbox will violate a placement restriction, such as a durability level, the primary head node selects extents outside of the sand box.

[0130] FIG. 13 is a high-level flowchart illustrating head nodes of a data storage unit performing a fail over operation in response to a failure of or loss of communication with one of the head nodes of the data storage unit, according to some embodiments.

[0131] At 1302 communication with a primary head node is lost or the primary head node fails. In some embodiments, a client device may lose contact with a primary head node and the client device may contact the secondary head node. This may trigger the secondary head node to attempt to take over as primary head node.

[0132] At 1304, in response to the secondary head node attempting to take over as primary head node, the local control plane issues a new sequence number to the secondary head node. The new sequence number may be greater than a sequence number previously issued to the previous primary head node. The new sequence number may be used by the secondary head node to gain write access to extents that were previously reserved for write access only by the previous primary head node.

[0133] At 1306, the secondary head node assumes the role of primary head node and begins to service writes directed to the volume. In some embodiments, the secondary head node may assume the role of primary head node by presenting the new sequence number received from the local control plane to sled controllers of the data storage system and receiving, from the sled controllers, credentials for writing to columns that store data of the volume.

[0134] At 1308, the local control plane designates another head node of the data storage unit to function as a secondary

head node for the volume or volume partition. Note that the previous secondary head node has assumed the role of primary head node, such that the volume is without a secondary head node causing the local control plane to designate a new secondary head node.

[0135] At 1310, the new primary head node (previous secondary head node) replicates log and index data for the volume to the newly designated secondary head node. In some embodiments, replicating log and index data may include replicating index data for the volume including pointers for volume data stored in data storage sleds of a data storage unit and volume data stored in the log of the new primary head node (previous secondary head node) that has not yet been flushed to the data storage sleds.

[0136] FIG. 14 is a block diagram illustrating performance and/or usage metrics being collected and accumulated in a data storage unit, according to some embodiments.

[0137] Data storage system 1400 may be the same as data storage system 900 illustrated in FIG. 9. Data storage system includes zonal control plane 1404 and data storage units 1406, 1428, and 1430. In some embodiments, data storage sleds and head nodes of a data storage unit may report performance and usage metrics to a local control plane for the data storage unit. For example, head nodes 1408, 1412, and 1414 of data storage unit 1406 report performance and usage metrics to local control plane 1410 of data storage unit 1406. Also, a sled controller of each of data storage sleds 1416 may report performance and usage metrics to local control plane 1410. In a similar manner, data storage sleds 1426 and head nodes 1418, 1420, and 1424 of data storage unit 1428 may report performance and usage metrics to local control plane 1422 of data storage unit 1428. Likewise, data storage sleds 1440 and head nodes 1432, 1434, and 1436 of data storage unit 1430 may report performance and usage metrics to local control plane 1438. In some embodiments, each local control plane of a data storage unit may in turn report accumulated performance and usage metrics to a zonal control plane for the data storage system. For example, local control planes 1410, 1422, and 1438 report performance and usage metrics to zonal control plane 1404. In some embodiments local control planes may use performance and usage metrics to balance loads between head nodes and to update sandbox recommendations that indicate recommended data storage sleds from which head nodes should select extents for a given volume. Also, a zonal control plane may use cumulative performance and usage metrics to balance volume assignments and/or move volumes between data storage units. In some embodiments, performance and usage metrics may be used by a local control plane to balance loads within a given data storage unit and accumulated performance and usage metrics may be used by a zonal control plane to balance loads between data storage units.

Input/Output Fencing of Mass Storage Devices from Unauthorized Head Nodes

[0138] In some embodiments, a sled controller of a data storage sled may implement a fencing protocol that prevents unauthorized head nodes from writing data to columns of mass storage devices located in a data storage sled along with the sled controller. In some embodiments, a sled controller may issue credentials or tokens to head nodes for accessing columns allocated to a particular volume serviced by the respective head nodes. The sled controller may only issue a new token to a head node if a column associated with

the credential or token is not currently reserved or if a head node seeking to access the column presents a sequence number greater than a sequence number stored for the column that indicates a sequence number of a previous head node that requested to access the column. For example, a newly designated primary head node for a given volume may receive from a local or zonal control plane a sequence number for the given volume that is greater than a previously issued sequence number for the given volume. The newly designated primary head node may then present the new sequence number to sled controllers of data storage sleds that include columns allocated for the volume. The sequence number of the newly designated primary head node may be greater than a sequence number stored in the columns that corresponded to a sequence number of a previous primary head node that accessed the columns. Upon determining that the newly designated primary head node has presented a sequence number greater than a stored sequence number, the sled controllers may issue a new token to the newly designated primary head node for accessing the columns.

[0139] For example, FIG. 15 illustrates interactions between a local control plane, head nodes, and data storage sleds of a data storage unit in relation to writing data to mass storage devices of a data storage sled of the data storage unit, according to some embodiments. Various interactions are illustrated between local control plane 1502 of data storage unit, head nodes 1504 and 1506 of the data storage unit and sled controllers 1508 of the data storage unit. Any of the data storage units described herein may include a local control plane, head nodes and sled controllers of data storage sleds that function as described in FIG. 15.

[0140] Phases 1, 2, and 3 are illustrated to show interactions that take place at different phases of operation of a data storage system. For example, phase 1 may be a normal phase in which a head node is assuming the role of primary head node for a volume or volume partition and functioning as the primary head node for the volume or volume partition. Phase 2 may represent a failover phase in which a secondary head node is assuming the role of primary head node for the volume, and phase 3 may represent a new normal phase wherein a newly designated primary head node is functioning as a primary head node for the volume.

[0141] At phase 1, local control plane 1502 assigns (1510) head node 1504 to be a primary head node for a volume and assigns (1512) head node 1506 to be a secondary head node for the volume. Assignment 1510 may include a new sequence number that is a monotonically increasing number that is greater than all sequence numbers previously issued by the local control plane 1502. At phase 1, in order to reserve columns of mass storage devices in different ones of multiple data storage sleds of a data storage unit, head node 1504 presents (1514) the new sequence number to sled controllers 1508 and reserves (1514) columns on mass storage devices located in data storage sleds that include the sled controllers 1508. At 1516, the sled controllers issue credentials or tokens to head node 1504 indicating that the columns are reserved for the volume and that head node 1504 is functioning as primary head node for the volume. At 1518, head node 1504 then issues a write request to sled controllers 1508 and includes along with the write requests the tokens or credentials issued by the sled controllers. The sled controllers verify that the credentials or tokens included with the write request are valid, perform the requested write, and at 1520 issue a write acknowledgement to head node

1504. Also the sled controllers store the sequence number and volume ID or volume partition ID in each column along with the data included with the write request.

[0142] During phase 2 or the fail over phase, communication is lost with head node **1504** at **1522**. In some embodiments, loss of communication with a primary head node may be triggered by a client device failing to be able to reach the primary head node and instead contacting the secondary head node. In such embodiments, the secondary head node may attempt to take over as primary head node (not illustrated in FIG. 15). In some embodiments, a local control plane may determine that a primary head node has been lost. In response to determining that a primary head node has failed or there is a loss of communication with a primary head node, at **1524**, local control plane **1502** promotes head node **1506** to primary head node for the volume and issues a new sequence number to head node **1506**. Head node **1506** then, at **1526**, presents the new sequence number issued to head node **1506** to sled controllers **1508** and requests access to the columns that store data for the volume for which head node **1506** is now the primary head node. The new sequence number issued to head node **1506** is greater than the sequence number issued to head node **1504** at **1510**. At **1528**, the sled controllers issue a new token or credential to head node **1506** that supersedes the token or credential issued to head node **1504** at **1516**.

[0143] During phase 3, head node **1506** functions as a primary head node for the volume. At **1530** head node **1506** includes with subsequent write requests the tokens issued from the sled controllers at **1528**. At **1532** sled controllers acknowledge subsequent writes from head node **1506**. Also, at **1534** head node **1504** that has lost communication with control plane **1502** and/or head node **1506** attempts to perform a write to columns assigned to the volume. However, subsequent to the failover, head node **1504** is no longer the primary head node for the volume and head node **1506** is functioning as primary head node for the volume. Thus, head node **1506** has exclusive access to columns of mass storage devices of extents allocated to the volume. Thus, at **1534** when head node **1504** attempts to access the columns sled controllers **1508** decline (**1536**) to perform the write. In addition, at **1536** the head node **1504** may read the volume ID and new sequence number stored in the columns assigned to the volume. The columns may store the new sequence number issued to head node **1506** during the failover. Upon determining that a new sequence number has been stored that supersedes the sequence number last issued to head node **1504**, head node **1504** may determine that it is no longer primary head node for the volume and may assume a role of secondary head node for the volume.

[0144] Note that each column stores a volume or volume partition ID for a volume for which the column is allocated along with a most recent sequence number. The volume ID and sequence number may be saved in persistent memory of the column. Also, a sled controller may store volume ID and sequence number information in a volatile memory of the sled controller. However, when a sled controller is reset, e.g. loses power, the volume and sequence number stored in the sled controller may be lost. However, volume and sequence number information stored in columns of mass storage devices may be persisted. This avoids complications that may arise if mass storage devices are moved between data storage sleds. For example, if a mass storage device is moved within a data storage sled or amongst data storage

sleds, sled controller volume ID and sequence number information may become inaccurate. However, because volume ID and sequence number information is lost from a sled controller whenever power is lost to the sled controller, the sled controller may be reset when a sled is removed from a data storage unit to access mass storage devices in the data storage sled avoiding such complications. Thus, subsequent to a reboot of a sled controller, head nodes serving as primary head nodes for volumes that have columns allocated on a sled that includes the sled controller may need to reclaim the columns. For example the head nodes may present respective sequence numbers issued to the head nodes and the sled controllers may issue new credentials or tokens to the head nodes if the sequence numbers presented have not be superseded, e.g. the sequence numbers stored in the columns are not greater than the sequence numbers being presented by the head nodes.

[0145] FIG. 16 is a high-level flowchart of a head node of a data storage unit flushing data stored in a storage of the head node to a data storage sled of the data storage unit, according to some embodiments.

[0146] At **1602**, a head node functioning as a primary head node for a volume receives a write request. At **1604**, the head node writes data included with the write request to a storage of the head node, such as a log of the head node.

[0147] At **1606**, in response to determining data stored in the storage of the head node exceeds a threshold, the head node requests sled controllers of multiple data storage sleds cause portions of the data stored in the storage of the head node be stored in multiple portions of different mass storage devices in different ones of the data storage sleds of the data storage unit. Requesting the sled controllers to store the data may further include presenting credentials (**1608**), such as credentials described in FIG. 15, to each of the sled controllers.

[0148] FIG. 17 is a high-level flowchart of a sled controller of a data storage sled processing a write request, according to some embodiments.

[0149] At **1702**, a sled controller receives a credential from a head node along with a write request. At **1704** and **1706**, the sled controller determines if the credential received at **1702** is a currently valid credential for a column of a mass storage device in a data storage sled that includes the sled controller. A sled controller may compare a sequence number and/or volume ID included in the credential with a sequence number and/or volume ID saved in the column for which access is requested. If the sequence number and/or volume ID included in the credential match the sequence number and/or volume ID stored in the column the sled controller may determine that the credential is valid. In some embodiments, a sled controller may store information that corresponds with a token or credential, such as a token number. If the information that corresponds with the token stored by the sled controller matches information included in the token, the sled controller may determine the credential or token is a valid credential. If a sequence number included in the credential or token is inferior to a sequence number stored in the column, the sled controller may determine that the credential or token is invalid. In some embodiments, a head node may not currently have credentials for a particular column and may present a sequence number that is greater than a stored sequence number stored for the column and the sled controller may issue credentials that supersede all previously issued cre-

credentials for the column, such as a new token that supersedes all tokens previously issued for the column.

[0150] At 1712, in response to determining at 1706 that the credential included with the write request is an invalid credential, the sled controller does not perform the requested write and returns a message to the head node indicating that the credential is invalid.

[0151] At 1708, in response to determining the credential is valid, the sled controller performs the requested write to the requested column of a mass storage device in the data storage sled along with the sled controller. At 1710 the sled controller acknowledges the write has been performed to the head node.

Data Storage Unit Design with Redundant Networks and Redundant Power

[0152] In some embodiments, a data storage unit may include redundant network and redundant power supplies and power distribution systems. Such redundant systems may reduce probabilities of failure thus allowing, for example, a single rack to store all parts of a volume while still meeting customer requirements for reliability and data durability. However, in some embodiments, a volume or volume partition may be stored in more than one data storage unit.

[0153] FIGS. 18A-D illustrate a data storage unit with redundant network paths within the data storage unit, according to some embodiments. Data storage unit 1850 illustrated in FIGS. 18A-D may be the same as data storage unit 100 illustrated in FIG. 1, or any other data storage unit described herein. FIGS. 18A-D further illustrate communication paths between network switches 1802 and 1804, head nodes 1806-1808, and data storage sleds 1834-1844. As can be seen, at least two redundant networks, including internal network 1852 and internal network 1854, are implemented within data storage unit 1850. Note that paths between components of data storage unit 1850 are illustrated on either side of data storage unit 1850 for clarity, but in practice paths between components of a data storage unit may be within the data storage unit over wires, cables, busways, etc. of the data storage unit.

[0154] In FIG. 18A redundant communication paths are established between head node 1806 and network 1828 via network switches 1802 and 1804. In some embodiments, a head node, such as head node 1806, may be assigned redundant network addresses routable from devices external to data storage unit 1850, such as public IP addresses, and may be reachable via either one of the network address using either one of network switches 1802 and 1804.

[0155] FIG. 18B illustrates redundant network paths between head nodes. For example head node 1806 may reach head node 1808 via internal network 1852 or internal network 1854, wherein internal network 1852 is via network switch 1802 and internal network 1854 is via network switch 1804. Note that there is a single network hop between head node 1806 and head node 1808 via network switch 1802 or network switch 1804. In some embodiments, a data storage unit may have a single network hop between head nodes and data storage sleds so that input/output operations do not require multiple network hops to retrieve or write data, thus improving IOPS performance and latency.

[0156] FIG. 18C illustrates redundant network paths between head nodes and data storage sleds. For example, head node 1806 may reach any of data storage sleds 1834-1844 via sled controllers 1812-1822. Each sled controller

may include two network ports that each are connected to different ones of internal networks 1852 or 1854 via either one or network switches 1802 or 1804. In some embodiments, each head node may be assigned at least two private network addresses and each sled controller may be assigned at least two private network addresses. The private network addresses assigned to the head nodes and sled controllers of the data storage sleds may enable the head nodes and sled controllers to communicate with each other via either one of internal networks 1852 or 1854. FIG. 18D illustrates a head node sending a response communication to a client device via either one of internal networks 1852 or 1854.

[0157] In some embodiments, a data storage unit may be configured to accept more or less head nodes in a rack of the data storage unit or to accept more or less data storage sleds in the rack of the data storage unit. Thus, compute capacity and data storage capacity of a data storage unit may be adjusted by varying a quantity of head nodes and/or data storage sleds that are included in the data storage unit.

[0158] FIGS. 19A-C illustrate a data storage unit configured to allow scaling of storage capacity and processing capacity, according to some embodiments. For example, data storage unit 1902 is shown in arrangement 1900 in FIG. 19A and in arrangement 1920 in FIG. 19B. In arrangement 1920 data storage unit 1902 includes more data storage sleds than in arrangement 1900. Also, in arrangement 1940, data storage unit 1902 includes more head nodes than in arrangement 1900. In some embodiments, a ratio of head nodes to data storage sleds may be adjusted to meet customer needs.

Example Computer System

[0159] FIG. 20 is a block diagram illustrating an example computer system, according to various embodiments. For example, computer system 2000 may be configured to implement storage and/or head nodes of a data storage unit, storage and/or a sled controller of a data storage sled, other data stores, and/or a client, in different embodiments. Computer system 2000 may be any of various types of devices, including, but not limited to, a personal computer system, desktop computer, laptop or notebook computer, mainframe computer system, handheld computer, workstation, network computer, a consumer device, application server, storage device, telephone, mobile telephone, or in general any type of computing device.

[0160] Computer system 2000 includes one or more processors 2010 (any of which may include multiple cores, which may be single or multi-threaded) coupled to a system memory 2020 via an input/output (I/O) interface 2030. Computer system 2000 further includes a network interface 2040 coupled to I/O interface 2030. In various embodiments, computer system 2000 may be a uniprocessor system including one processor 2010, or a multiprocessor system including several processors 2010 (e.g., two, four, eight, or another suitable number). Processors 2010 may be any suitable processors capable of executing instructions. For example, in various embodiments, processors 2010 may be general-purpose or embedded processors implementing any of a variety of instruction set architectures (ISAs), such as the x86, PowerPC, SPARC, or MIPS ISAs, or any other suitable ISA. In multiprocessor systems, each of processors 2010 may commonly, but not necessarily, implement the same ISA. The computer system 2000 also includes one or more network communication devices (e.g., network inter-

face 2040) for communicating with other systems and/or components over a communications network (e.g. Internet, LAN, etc.).

[0161] In the illustrated embodiment, computer system 2000 also includes one or more persistent storage devices 2060 and/or one or more I/O devices 2080. In various embodiments, persistent storage devices 2060 may correspond to disk drives, tape drives, solid state memory, other mass storage devices, block-based storage devices, or any other persistent storage device. Computer system 2000 (or a distributed application or operating system operating thereon) may store instructions and/or data in persistent storage devices 2060, as desired, and may retrieve the stored instruction and/or data as needed. For example, in some embodiments, computer system 2000 may host a storage unit head node, and persistent storage 2060 may include the SSDs that include extents allocated to that head node.

[0162] Computer system 2000 includes one or more system memories 2020 that are configured to store instructions and data accessible by processor(s) 2010. In various embodiments, system memories 2020 may be implemented using any suitable memory technology, (e.g., one or more of cache, static random access memory (SRAM), DRAM, RDRAM, EDO RAM, DDR 10 RAM, synchronous dynamic RAM (SDRAM), Rambus RAM, EEPROM, non-volatile/Flash-type memory, or any other type of memory). System memory 2020 may contain program instructions 2025 that are executable by processor(s) 2010 to implement the methods and techniques described herein. In various embodiments, program instructions 2025 may be encoded in platform native binary, any interpreted language such as Java™ byte-code, or in any other language such as C/C++, Java™, etc., or in any combination thereof. For example, in the illustrated embodiment, program instructions 2025 include program instructions executable to implement the functionality of a storage node, in different embodiments. In some embodiments, program instructions 2025 may implement multiple separate clients, nodes, and/or other components.

[0163] In some embodiments, program instructions 2025 may include instructions executable to implement an operating system (not shown), which may be any of various operating systems, such as UNIX, LINUX, Solaris™, MacOS™, Windows™, etc. Any or all of program instructions 2025 may be provided as a computer program product, or software, that may include a non-transitory computer-readable storage medium having stored thereon instructions, which may be used to program a computer system (or other electronic devices) to perform a process according to various embodiments. A non-transitory computer-readable storage medium may include any mechanism for storing information in a form (e.g., software, processing application) readable by a machine (e.g., a computer). Generally speaking, a non-transitory computer-accessible medium may include computer-readable storage media or memory media such as magnetic or optical media, e.g., disk or DVD/CD-ROM coupled to computer system 2000 via I/O interface 2030. A non-transitory computer-readable storage medium may also include any volatile or non-volatile media such as RAM (e.g. SDRAM, DDR SDRAM, RDRAM, SRAM, etc.), ROM, etc., that may be included in some embodiments of computer system 2000 as system memory 2020 or another type of memory. In other embodiments, program instructions may be communicated using optical, acoustical or

other form of propagated signal (e.g., carrier waves, infrared signals, digital signals, etc.) conveyed via a communication medium such as a network and/or a wireless link, such as may be implemented via network interface 2040.

[0164] In some embodiments, system memory 2020 may include data store 2045, which may be configured as described herein. In general, system memory 2020 (e.g., data store 2045 within system memory 2020), persistent storage 2060, and/or remote storage 2070 may store data blocks, replicas of data blocks, metadata associated with data blocks and/or their state, configuration information, and/or any other information usable in implementing the methods and techniques described herein.

[0165] In one embodiment, I/O interface 2030 may be configured to coordinate I/O traffic between processor 2010, system memory 2020 and any peripheral devices in the system, including through network interface 2040 or other peripheral interfaces. In some embodiments, I/O interface 2030 may perform any necessary protocol, timing or other data transformations to convert data signals from one component (e.g., system memory 2020) into a format suitable for use by another component (e.g., processor 2010). In some embodiments, I/O interface 2030 may include support for devices attached through various types of peripheral buses, such as a variant of the Peripheral Component Interconnect (PCI) bus standard or the Universal Serial Bus (USB) standard, for example. In some embodiments, the function of I/O interface 2030 may be split into two or more separate components, such as a north bridge and a south bridge, for example. Also, in some embodiments, some or all of the functionality of I/O interface 2030, such as an interface to system memory 2020, may be incorporated directly into processor 2010.

[0166] Network interface 2040 may be configured to allow data to be exchanged between computer system 2000 and other devices attached to a network, such as other computer systems 2090, for example. In addition, network interface 2040 may be configured to allow communication between computer system 2000 and various I/O devices 2050 and/or remote storage 2070. Input/output devices 2050 may, in some embodiments, include one or more display terminals, keyboards, keypads, touchpads, scanning devices, voice or optical recognition devices, or any other devices suitable for entering or retrieving data by one or more computer systems 2000. Multiple input/output devices 2050 may be present in computer system 2000 or may be distributed on various nodes of a distributed system that includes computer system 2000. In some embodiments, similar input/output devices may be separate from computer system 2000 and may interact with one or more nodes of a distributed system that includes computer system 2000 through a wired or wireless connection, such as over network interface 2040. Network interface 2040 may commonly support one or more wireless networking protocols (e.g., Wi-Fi/IEEE 802.11, or another wireless networking standard). However, in various embodiments, network interface 2040 may support communication via any suitable wired or wireless general data networks, such as other types of Ethernet networks, for example. Additionally, network interface 2040 may support communication via telecommunications/telephony networks such as analog voice networks or digital fiber communications networks, via storage area networks such as Ethernet, Fibre Channel SANs, or via any other suitable type of network and/or protocol. In various embodiments, computer system

2000 may include more, fewer, or different components than those illustrated in FIG. 20 (e.g., displays, video cards, audio cards, peripheral devices, other network interfaces such as an ATM interface, an Ethernet interface, a Frame Relay interface, etc.)

[0167] It is noted that any of the distributed system embodiments described herein, or any of their components, may be implemented as one or more network-based services. For example, a compute cluster within a computing service may present computing and/or storage services and/or other types of services that employ the distributed computing systems described herein to clients as network-based services. In some embodiments, a network-based service may be implemented by a software and/or hardware system designed to support interoperable machine-to-machine interaction over a network. A network-based service may have an interface described in a machine-processable format, such as the Web Services Description Language (WSDL). Other systems may interact with the network-based service in a manner prescribed by the description of the network-based service's interface. For example, the network-based service may define various operations that other systems may invoke, and may define a particular application programming interface (API) to which other systems may be expected to conform when requesting the various operations. though

[0168] In various embodiments, a network-based service may be requested or invoked through the use of a message that includes parameters and/or data associated with the network-based services request. Such a message may be formatted according to a particular markup language such as Extensible Markup Language (XML), and/or may be encapsulated using a protocol such as Simple Object Access Protocol (SOAP). To perform a network-based services request, a network-based services client may assemble a message including the request and convey the message to an addressable endpoint (e.g., a Uniform Resource Locator (URL)) corresponding to the network-based service, using an Internet-based application layer transfer protocol such as Hypertext Transfer Protocol (HTTP).

[0169] In some embodiments, network-based services may be implemented using Representational State Transfer ("RESTful") techniques rather than message-based techniques. For example, a network-based service implemented according to a RESTful technique may be invoked through parameters included within an HTTP method such as PUT, GET, or DELETE, rather than encapsulated within a SOAP message.

[0170] Although the embodiments above have been described in considerable detail, numerous variations and modifications may be made as would become apparent to those skilled in the art once the above disclosure is fully appreciated. It is intended that the following claims be interpreted to embrace all such modifications and changes and, accordingly, the above description to be regarded in an illustrative rather than a restrictive sense.

What is claimed is:

1. A data storage system comprising:
 - a plurality of head nodes; and
 - a plurality of data storage sleds,

wherein for a partition of a volume to be stored in the data storage system, a particular one of the head nodes is designated as a primary head node for the volume partition and

another one of the head nodes is designated as a secondary head node for the volume partition,

wherein, based, at least in part, on receiving a write request for the volume partition, the head node designated as the primary head node for the volume partition is configured to:

- write data included with the write request to a storage of the head node designated as the primary head node; and
- cause the data included with the write request to be replicated to the other head node designated as the secondary head node;

wherein the head node designated as the primary head node for the volume partition is further configured to cause respective parts of the data stored in the storage of the head node to be erasure encoded and stored in a plurality of respective mass storage devices.

2. The data storage system of claim 1, wherein the plurality of mass storage devices that store the erasure encoded respective parts of the data are in different ones of the plurality of data storage sleds of the data storage system.

3. The data storage system of claim 1, wherein at least a portion of the data written to the storage of the head node designated as the primary head node is caused to be erasure encoded and stored in the plurality of respective mass storage devices based, at least in part, on:

- a quantity of data stored in the storage of the head node designated as the primary head node exceeding a threshold amount of stored data; or

- an amount of time, since data stored in the storage of the head node designated as the primary head node was stored, exceeding a threshold amount of time.

4. The data storage system of claim 1, wherein the head node designated as the primary head node is configured to release storage space of the head node subsequent to performing said cause respective parts of the data stored in the storage of the head node to be stored in the plurality of respective mass storage devices; and

- wherein the head node designated as the primary head node is configured to cause storage space of the other head node designated as the secondary head node to be released subsequent to performing said cause respective parts of the data stored in the storage of the head node to be stored in the plurality of respective mass storage devices.

5. A data storage system comprising:

- a head node of the data storage system;

wherein, based, at least in part, on receiving a write request for a volume partition, the head node acting as a primary head node of the data storage system for the volume partition is configured to:

- write data included with the write request to a storage of the head node; and

- cause the data included with the write request to be replicated to another head node of the data storage system designated as a secondary head node for the volume partition;

wherein the head node is further configured to cause respective parts of the data stored in the storage of the head node to be stored in a plurality of respective mass storage devices in different ones of a plurality of data storage sleds of the data storage system.

6. The data storage system of claim 5, wherein the data is stored at a particular level of data durability when stored in

the storage of the head node acting as primary head node and the other the other head node designated as secondary head node; and

wherein the data is stored at another level of data durability when the respective parts of the data are stored in the plurality of respective mass storage devices each in different ones of the plurality of data storage sleds of the data storage system.

7. The data storage system of claim 6, wherein the head node designated as the primary head node is configured to cause the data included in the write request to be replicated to one or more additional head nodes.

8. The data storage system of claim 6, wherein the head node designated as the primary head node is configured to adjust a number of data storage sleds across which the data stored in the storage of the head node is caused to be stored.

9. The data storage system of claim 5, wherein the storage of the head node and the storage of the other head node are log-structured storages.

10. The data storage system of claim 9, wherein the storage of the head node comprises an index, wherein for the volume partition, the index identifies a location within the log of the head node at which the data included in the write request is stored,

wherein the head node is configured to update the index to identify respective portions of the plurality of respective mass storage devices each in different ones of the plurality of data storage sleds where the respective parts of the data are stored.

11. The data storage system of claim 5, wherein the head node is configured to cause respective parts of the data stored in the storage of the head node to be stored in a plurality of respective mass storage devices asynchronously with the data being stored in the storage of the head node and in a storage of the head node designated as the secondary head node.

12. The data storage system of claim 5, wherein to cause the respective parts of the data stored in the in the storage of the head node to be stored in the plurality of respective mass storage devices, the head node is further configured to erasure encode the data stored in the storage of the head node across the plurality of mass storage devices.

13. The data storage system of claim 5, wherein the head node is configured to be a secondary head node for another volume partition stored in the data storage system.

14. The data storage system of claim 13, wherein the head node designated as a secondary head node for the other volume partition does not cause data of the other volume partition to be stored in mass storage devices of the data storage sleds; and

wherein the head node, upon being designated as a primary head node for the other volume partition is configured to cause respective parts data of the other volume partition stored in the storage of the head node to be stored in a plurality of respective mass storage devices.

15. A non-transitory computer readable medium storing program instructions for implementing a data storage sys-

tem, wherein the program instructions when executed by a processor cause the data storage system to:

receive a write request for a volume partition:

write data included with the write request to a storage of a head node designated as primary head node for the volume partition; and

replicate the data included in the write request to another head node of the data storage system designated as a secondary head node for the volume partition;

erasure encode at least a portion of the data stored in the storage of the head node designated as the primary head node; and

store respective parts of the erasure encoded data in a plurality of respective mass storage devices each in different ones of a plurality of data storage sleds of the data storage system.

16. The non-transitory computer readable medium of claim 15, wherein the program instructions when executed cause the data storage system to asynchronously perform said:

replicate the data to another head node of the data storage system designated as a secondary head node for the volume partition; and

erasure encode the at least a portion of the data and store the respective parts of the erasure encoded data.

17. The non-transitory computer readable medium of claim 16, wherein the program instructions when executed cause the data storage system to:

cause respective parts of the data stored in the storage of the head node designated as the primary head node to be stored in a plurality of respective mass storage devices, based, at least in part, on:

an amount of time since the data was written to the storage of the head node designated as primary head node; or an amount of data stored in the storage of the head node designated as the primary head node exceeding a threshold amount of data.

18. The non-transitory computer readable medium of claim 15, wherein the program instructions when executed cause the data storage system to:

erasure encode the data stored in the storage of the head node designated as the primary head node across the plurality of mass storage devices.

19. The non-transitory computer readable medium of claim 15, wherein the erasure coded data is stored across at least six mass storage devices, wherein at least four of the mass storage devices store portions of the data and at least two of the mass storage devices store coded data derived from the data.

20. The non-transitory computer readable medium of claim 15, wherein the program instructions when executed cause the data storage system to:

detect a failure of one or more of the mass storage devices in a particular one of the data storage sleds; and

copy data stored in remaining mass storage devices of the data storage sled to mass storage devices in one or more other data storage sleds.

* * * * *